

HP

MAY 2000

Professional

Servers


Workstations

Services

Applications

www.hppro.com

BUSINESS INTELLIGENCE



Information Integration

Data Management:
Playing It Smart PAGE 12

Data Mining Solutions PAGE 18

META Group's Data Warehouse Scorecard PAGE 21

An Interview with Bill Inmon PAGE 23

HP e3000 PAGE 26

HP-UX Admin Man
Fault Management

On the Server Side
Updating Your Microsoft Certification

Product Watch
Bluestone's XML for HP-UX

#BXNDXGF*****ALL FOR ADC 773
#17810/CH/4# H HP003 SCK 619
JERRY FOCHTMAN, DIR HP3000 PROD DEVELOPMENT
BRADMARK TECHNOLOGIES
2818 MOUNTAIN GREEN TRL
HUMBLE TX 77345-2215


NO MATTER WHAT
HAPPENS

NO MATTER WHERE
IT HAPPENS


WE KEEP IT

FROM HAPPENING

TO YOU



© 1999 Hewlett-Packard Co.



From something as simple as a power outage, or as unstoppable as a flood, HP Business Recovery Services can ensure your business's continuity. We are a globally-capable recovery company that's been successfully delivering disaster assistance around the world for over 12 years. Our service mission is to meet recovery time objectives – the time lapse between disaster declaration and reconnection to critical applications – from less than one hour to 72 hours. Call us at 800-863-5360, or visit us at www.hp.com/go/recovery. We'll tell you more about how we can protect you and keep you running. So, no matter what, or no matter where, we'll get you back to business.



MAY 2000

Vol. 15, No. 5

HP Professional

BUSINESS INTELLIGENCE

12 Smart Ways of Managing Data: Data Warehousing and Business Intelligence

Data. It is accumulating so rapidly, and in so many different sources, that organizations are struggling to consolidate the information and provide access to it across the enterprise and within the supply chain. UNIX remains the number one platform for data warehousing, and HP-UX servers provide the scalability needed to run BI/DW solutions. **By Jean Nattkemper**

DATA MINING

18 How to Mine Your Own Business

Data mining technology takes information about how the elements in a warehouse are related and uses technology grounded in statistics and neural networks to look for patterns between values that could be significant. Automating the ability to pick up on these potentially revealing patterns can prove valuable, but it places serious requirements on the skills of managers and warehouse architecture design. **By Katherine Hammer**

DATA WAREHOUSING

21 Data Warehouse Factors to Address for Success

The responses in META Group's recent data warehouse survey confirm that team experience, architecture, expectation management and data quality are the prevailing factors influencing the success of DW initiatives. With the emerging knowledge base of DW best practices and a growing skill base of DW practitioners, over 90 percent of DW installations are successful. **By Doug Laney**

23 A One-on-One Interview with Bill Inmon

The "Father of the Data Warehouse" discusses what's in store for the industry. **By Charlie Simpson**


HP e3000

26 ASAP Software Integrates E-Commerce Front End to Its HP e3000 Business Systems

Based on company growth and the need for greater extensibility and user friendliness, ASAP Software needed to replace a homegrown package to keep pace with a growing e-commerce business. HP Partner, Fioravanti-Redwood International had a solution. **By Bill Pike**

30 Get the Most out of Data on the HP e3000

When HP's Commercial Systems Division (CSY) dropped an "e" in front of the HP 3000, they told the world that the platform is a player in e-commerce and e-services, as well as data warehousing and business intelligence. Find out if they really meant it. **By Jean Nattkemper**



Cover Illustration by Bob Daly

COLUMNS

4 Guest Editorial: The New World of Business Intelligence

Douglas Hackney, President of
Enterprise Group Ltd.

32 The HP-UX Admin Man: It's Not My Fault

Fred's been reading about "fault
management" and shares some of
the blame about what he's learned.

By Fred Mallett

34 On the Server Side: Updating Your Microsoft Certification

With Windows 2000's arrival, it's
time to update your professional cer-
tifications from Microsoft. Ryan looks
at the pros and cons. **By Ryan Maley**

PRODUCT WATCH

10 Bluestone's XML for HP-UX

DEPARTMENTS

6 The Buzz

36 Marketplace

39 Advertiser Index

40 New Products

www.hpprofessional.com

Everything HP

your work
your world
your technology

Now Available! HP 3000 Rentals XP 256

If it says HP—we deal in it!

P405 Dot Matrix Printer

\$2195

- RuggedWriter replacement
- Printer interface cable must be purchased separately
- This printer must be ordered with one of four support options. Call TSA for the pricing and warranty options
- Optional Automatic Sheet Feeders available
- All consumable items available

1-800-422-4872

COMPLETE LINE OF SERVERS AND WORKSTATIONS IN STOCK

• Order your N Class today...have it tomorrow!

• Y2K Special-K580's

• PC's, Hardcopy, new and refurbished



Everything HP

www.tsa.com • 713-935-1500
Fax: 713-935-1555

2040 West Sam Houston Parkway N. • Houston, Texas 77043

HP Professional

www.hppro.com

Senior VP & Publishing Director Thomas J. Wilson
twilson@hppro.com
Editorial Director Charlie Simpson
csimpson@hppro.com
Managing Editor Tom Schaffner
tschaffner@hppro.com
Editor at Large Jean Nattkemper
jnattkemper@hppro.com
Associate Editor/Web Editor Carolyn Majewski
cmajewski@hppro.com
Assistant Editor Traci Savidge
tsavidge@hppro.com

COLUMNISTS

Workstations Fred Mallett
frederm@famece.com
Servers Ryan Maley
ryan@maley.org
HP e3000 Bill Pike
pikecom@earthlink.net

Creative Director Emily Rich
Art Director Robert Chapman
Graphic Designer Scott Rovin

Production Director Jennifer Burrow
Production Manager Joleen Mathews
Production Coordinator Laurie Chadwick

Circulation Manager Peggy Nelson

101communications

CHAIRMAN & CEO
Curt Hessler
PRESIDENT & COO
Jeffrey S. Klein
FOUNDER & VICE CHAIRMAN
Bill Slapin
SENIOR VP & PUBLISHING DIRECTOR
Thomas J. Wilson
SENIOR VP, CONFERENCES & EVENTS
Steven Crofts
VP, FINANCE
Andrew D. Landis
VP, CIRCULATION & DIRECT MARKETING
Christine Oldenbrook
VP, MANUFACTURING & PUBLISHING OPERATIONS
Bradford C. Stauffer
VP, MARKETING & BUSINESS DEVELOPMENT
Mark Sande
VP, HUMAN RESOURCES
Michael J. Valenti

SUBSCRIPTION SERVICES

HP PROFESSIONAL ISSN 0896-145X is published monthly by 101communications LLC, 9121 Oakdale Ave., Suite 101, Chatsworth California, 91311. Subscriptions are complimentary for qualified U.S. and Canadian sites. Periodicals postage paid at Chatsworth, CA 91311, and additional mailing offices. For address changes and other subscription information in the U.S.: hp@omeda.com or call (800) 306-6332. Outside the U.S., call (847) 291-5212 or fax (847) 564-9002. In the U.S., missed issues must be claimed within 45 days of the publication date; outside the U.S., missed issues must be claimed within 90 days. Editorial and advertising sales offices at 1300 Virginia Dr., Ste. 400, Fort Washington, PA 19034 • tel: (215) 643-8000, fax: (215) 643-3901.
US POSTMASTER: Send all correspondence and address changes to HP PROFESSIONAL, P.O. Box 3053, Northbrook, IL 60065. **COPYRIGHT** © 2000 by 101communications LLC. **CANADIAN POSTMASTER:** Send all correspondence and address changes to 101communications LLC, C/O N.L., P.O. Box 44, RPO Rockwood Mall, Mississauga, ON L4Z 9Z9.

PERMISSIONS

HP Professional, an independent magazine, is not affiliated with the Hewlett-Packard Company. HP and Hewlett-Packard are registered trademarks and HP Professional is a trademark of the Hewlett-Packard Company. Other trademarks and trade names used throughout the publication are the property of their respective owners.

All rights reserved. No part of this publication may be reproduced in any form without written permission from the publisher. All submitted manuscripts, photographs and/or artwork are sent to 101communications LLC at the sole risk of the sender. Neither 101communications LLC nor HP Professional magazine is responsible for any loss or damage.

Canada Post International Publication
Mail Canadian Distribution Sales Agreement #IPM0264431

Printed in U.S.A.

ABP



HP PROFESSIONAL EDITORIAL ADVISORY BOARD

Robert Bruen
Computing and Networking Manager at the MIT Lab for Nuclear Science

Thomas Kucharvy
President of Summit Strategies

John R. Logan
Vice President of Aberdeen Group Inc.

Michael J. Meinz
Principal Technical Consultant at General Mills Inc.

Bard F. White
CIO and Worldwide Director of MIS for Spalding Sports Worldwide

FOR REPRINTS OF ARTICLES IN ANY ISSUE, CONTACT

Reprint Management Services

Greenfield Corporate Center, 1811 Olde Homestead Lane
Lancaster, PA 17601

717.399.1900 • FAX 717.399.8900 • www.rmsreprints.com

Hewlett, Packard & Forsythe.

Seems we've always been close. Really close. Fact is, we know them as well as we know ourselves.

Guess it's one reason we're HP's leading e-solutions provider. It means a lot to us . . . both of us.

Could mean a lot to you too. Call 800-843-4488 or visit us at www.forsythesolutions.com.




The New World of Business Intelligence

If asked to identify a market that will be worth \$148 billion by 2003, you would probably think first of e-business, customer relationship management (CRM), supply chain management (SCM) or enterprise resource planning (ERP). Most people consider these market segments to be tomorrow's big guns when it comes to market growth and overall size.

Taken together, e-business, CRM, SCM and ERP get the lion's share of cover stories, feature articles and headlines in the trade press, as well as regular coverage in general-interest business publications, such as the *Wall Street Journal*. Venture capitalists are falling all over themselves to throw money at 20-year-old college students with loosely formed ideas for any of these market segments. In fact, startups can put the letter "e" in front of just about anything, including dog food, and attract millions of VC dollars.

Another market that is capturing far less attention; is growing at exponential rates; is mission-critical to every business, regardless of size; has every major technology company as a player; features e-everything delivery and interaction; offers exponentially more capability at orders of magnitude and lower prices than just a year ago; is commercializing top-secret technology from formerly off-limits U.S. government programs; boasts some of the hottest initial public offerings (IPO) and best-performing companies of the last two years; and has nowhere to go but up. That market is business intelligence (BI). Survey.com predicts that it will be a 148-billion-dollar global market by 2003.



BI is the sleeper market segment in the technology world. It is often the poor stepchild of the high-flying, big-headline e-business and CRM segments. Largely shunned through the '90s by the general press, and misunderstood by most technology and financial analysts, it has slowly been building strength and momentum for the last decade.

The current BI market is built on the foundation of a modern BI infrastructure, consisting of a federated architecture, accommodating all the components of a contemporary BI system: packaged/turnkey data warehouses and data marts, packaged/turnkey analytical applications, custom-built data warehouses and data marts, custom-built analytic applications, data mining, online analytical processing (OLAP) tools, query and reporting tools, production reporting tools, data quality tools, extraction transformation and load tools, system management tools, information delivery tools, enterprise information portals, reporting systems, knowledge management systems, database systems, etc. The BI federated architecture is the "big tent" that provides the foundation and environment to facilitate and enable business information flow, analysis and decision-making.

BI's evolution from a small-scale, rather provincial market to a powerful global segment has not come without some growing pains. In the early days, the BI market was defined and led by a handful of people, tools, technologies and companies. Some market leaders, such as Prism Solutions and Carleton, have fallen by the wayside, unable to keep up with the rapid pace of innovation and the challenges of today's extremely diverse and very fast-moving market. Some of yesterday's market leaders have been crushed by new, innovative approaches (bottom-up), architectures (federated systems), and technologies (second-generation extraction transformation and load tools) that have redefined the entire BI market. This trend will continue and promises to accelerate in the coming years.

Today, the BI market balances on the cusp between the old world and the new. The market is hitting the elbow in the hockey stick of the growth curve and is rapidly accelerating. Yesterday's BI world was a market of thousands. BI was a small market, formed and led by like-minded technologists comfortable in a rather insular small-scale world.

The BI market is moving quickly toward a market of millions. Rapidly dropping price points and powerful packaged/turnkey solutions – such as profitability, inventory, marketing campaign and sales analysis – are enabling businesses of all sizes and technical abilities to enjoy powerful BI capabilities. Business-driven innovators delivering pragmatic solutions are at the vanguard of the new world of business intelligence.

–Douglas Hackney, President
Enterprise Group Ltd.

Uptime Recipe #1


Get lots of servers and software.

Mix frantically.

Repeat indefinitely.

Uptime Recipe #2

Remove Stratus server from box.



If you want continuous availability for your HP-UX® operating environment, the choice is simple. Unless you're a chaos freak. Only Stratus® provides the highest level of fault tolerance without cumbersome software and service requirements. Our built-in redundancy means that if a component fails, its partner is already doing the same job — so there's not a moment of downtime. Plus, our customer support center is automatically alerted, diagnoses the problem and sends a replacement part. You may not even know there was a problem until the new part arrives. For 20 years, our servers have been performing in Fortune 500 companies where downtime is not an option. To see if you qualify for an on-site trial, visit us soon at www.stratus.com/uptime.

Stratus®

THE WORLD'S MOST RELIABLE SERVERS™

Stratus and the Stratus logo are registered trademarks of Stratus Computer, Inc. "The World's Most Reliable Servers" is a trademark of Stratus Computer Systems, S.a.r.l., Luxembourg. HP-UX is a registered trademark of Hewlett-Packard Company.



News and Trends for the HP Community

HP DELIVERS ON INFRASTRUCTURE INITIATIVES

Talk to almost anyone at HP and somewhere in the conversation, you'll hear, "We are in the infrastructure business." It's hard to doubt the commitment: HP continues to unveil initiatives and deals that position the company as a major provider of the infrastructure for e-services and the new e-world in general.

In two recent moves, HP announced Infrastructure-on-Tap, a new service that delivers back-end infrastructure to Web-centric businesses over the Internet, and a partnership with service provider, PRIMUS Telecommunication Group, to combine HP infrastructure with PRIMUS' services.

Infrastructure-on-Tap

Infrastructure-on-Tap, designed, owned and operated by HP, combines HP's UNIX and Windows NT systems, along with enabling software applications, and database management and infrastructure systems from HP and partners, like Cisco, Microsoft, Nortel Networks and Oracle. HP manages facilities and network services through alliances with wholesale service providers.

The service offers ASPs, hubs, portals and exchanges a scalable, high-availability infrastructure, but requires no upfront IT investment. "Leveraging the similarities in IT architecture requirements, HP

shifted the paradigm from 'you have to build it' to 'you have a choice,'" says Ann Livermore, President of HP's Enterprise and Commercial Business. "If you're a large enterprise, building your own IT infrastructure may make sense. But, most B2B hubs build only because they fear a sudden rise in business could pull down their Web sites and decrease revenues. Valuable resources are used for IT infrastructures that are not fully utilized – in case their capacity spikes periodically."

HP's initial targets for Infrastructure-on-Tap are business-to-business hubs, a lucrative market. IDC predicts that the number of B2B electronic marketplaces will grow from about 1,000 today to almost 10,000 within the next 18 months.

ERP vendor, Lawson Software, and IPNet Solutions, a provider of software used to create electronic trading exchanges, are HP's first customers for Infrastructure-on-Tap. Lawson is using the service to power its ijob electronic recruiting services.



HP Funds Start-Up to Deliver Infrastructure

HP has formed a deal with PRIMUS to supply the service provider with up to \$50 million, offered through a convertible debt arrangement. PRIMUS will use the financing to upgrade its global broadband network, designed for small to mid-size companies developing Web-based businesses.

The partnership combines HP infrastructure and PRIMUS' services. PRIMUS will deliver its Internet, data and voice services through HP-powered data centers deployed through the service provider's existing worldwide facilities-based communications network.

PRIMUS will spend much of the HP investment to expand its data centers in Western Europe, Australia, Japan, Brazil and other regions. The first of the 10,000-square-foot data centers based on HP technology is expected to open in the United Kingdom by June. PRIMUS intends to build another four data centers by the end of the year.

PRIMUS is deploying HP's UNIX and NT servers, storage products and software in the data centers, as well as relying on HP for technical support. Starting with the U.K. and Australia, and then expanding to other targeted countries, PRIMUS will also offer HP's Commerce for the Millennium solution on the PRIMUS global backbone network. Moreover, PRIMUS intends to enhance the HP solution by bundling other communications services with it.

Speedware Enters Wireless Application Development Space

Speedware Corporation, a provider of application development tools and business intelligence technology, has become one of the first tools vendors to break into the Internet wireless communications space. Internet wireless communications represents a burgeoning market, according to Forrester Research, which predicts that Americans will spend almost \$60 billion for wireless services in 2005.

Speedware's entries in the wireless arena are a graphical development tool designed specifically for developing wireless Internet applications, and new professional services aimed at helping companies build wireless applications.

The new tool, MobileDev, recognizes the differences between standard Internet development and development devised specifically for wireless Internet technology, according to Chris Koppe, Speedware's Director of Wireless Business Development. "There are no vendors creating tools specifically for wireless develop-

ment," Koppe says. "XML and Internet tools vendors support WML and HDML presentation layers, but these tools do not have the appropriate graphical representations for wireless objects, nor do they support key features that differentiate wireless applications from Internet applications."

The features Koppe refers to include subscriber ID, alert notifications and location-based services. Not currently supported by the Internet and Web browsers, they are key within the wireless space.

Speedware touts MobileDev for its "rich graphical interface," as well as its ability to speed development and provide for easy maintenance of wireless applications. "Currently, the development of applications for wireless is a slow process, but MobileDev allows for application development in record time – as much as five and 10 times over traditional methods," Koppe says.

ORACLE DATABASE SERVER

The industry is finally getting a look at the much-delayed Oracle8i Appliance, a database server previously code-named Raw Iron. The server appliance, now shipping through hardware OEMs and valued-added distributors, allows for deployment of the Oracle8i database independent of Windows NT or UNIX operating systems. Oracle is targeting the product at small and mid-size businesses.

The Oracle 8i Appliance features a database, processors, and parts of Sun's Solaris operating system. It also includes Oracle software that allows for remote management of the appliance over the Internet, and WebDB, a tool that enables businesses to develop and run their own Web-based software.

Approximately 50 percent of Oracle customers run their Oracle databases on dedicated servers, so it makes sense for Oracle to have pared down the operating system, limiting it to the functions required to run an Oracle database. With the server appliance, Oracle aims to reduce the cost of software administration, and enhance the performance of the database by allowing software operations to be more closely tuned for specific hardware.

Oracle's new offering has met with an enthusiastic response from hardware and software vendors. Hardware vendors signing up to ship the product include HP, Compaq, Dell and Fujitsu Siemens. Pioneer/Keylink and HallMark Global Solutions and have agreed to manufacture and distribute the product on non-branded Intel systems.


About 50 ISVs, including J.D. Edwards, Legato, Macromedia and TIBCO, also announced support for the new product. ISVs supporting the Oracle8i Appliance have validated their applications against the appliance at four newly created Application Solution Centers, established to test third-party applications with the database server.

COGNOS' BI PORTAL

Cognos unveiled their new EnterpriseService Portal (ESP), a customizable, personalized portal that allows Cognos business intelligence customers to receive the level of customer service, support and education they choose. Cognos will implement the portal in stages throughout the year.

The ESP provides access to training, including new Web-based training offers;

Nortel Leads Group to Develop Broadband Standards



Nortel Networks Corp. is heading up a group of more than 30 Internet infrastructure, content and service providers in an effort to set common standards for broadband delivery over the Internet. The group, called the Broadband Content Delivery Forum (BCDF), aims to drive broadband by making sure vendors develop common equipment and services, instead of fighting over competing technology.

Nortel has lined up some big-time support for its effort. BCDF boasts networking heavyweights HP, Sun and BT (British Telecommunications), as well as broadcasters like BBC (British Broadcasting Corporation) and NBC's Internet Unit as members.

Nortel has set some specific goals for the group, one of which is to accelerate delivery of data over the Internet. The group will also work to improve the quality of the transmission of high-speed video and data services. By doing this, BCDF plans to encourage use of these services and drive development of new products.

The group also plans to develop technology that allows networks to identify Web customers. Service providers could use this kind of technology to create personalized content and servers for their customers.

The first meeting of BCDF is scheduled for this month. Other members include Akamai Technologies, Aliant, Alta Vista, AT&T Broadband, Bertelsmann, Be Here, Broadband Digital Group, Broadjump, Canbras Communications, CMGI, DSL Networks, El Mundo, Enron, First Mark Comm, iKnowledge, InfoLibria, InfoSpace.com, Inktomi, NetActive, Qwest Communications International, Redo Ajato, Starguide Digital Network, Telocity, Telstra, TVA Sistema de Televisao, Universo Online and Zyan Communications.

technical support; and post-sales consulting through Cognos' Consulting Services organization. ESP delivers Web-based training to Cognos' customers anywhere in the world, at any time. It also provides access to live customer support, customer chats, FAQs, a solutions knowledge base, and other forms of knowledge transfer, such as online diagnostic support.

"Cognos is evolving from the brick-and-mortar customer service model to a new e-business model which is more in tune with the new e-economy," says Rob Ashe, Cognos Senior Vice President of Customer Service and Support. Ashe's comments about the "new e-economy" are underscored by a recent report by E-Offering, an online investment banking firm. The report suggests that the Internet and information technology are fundamentally changing the way companies conduct business, predicts that the

growth of business-to-business e-commerce will far exceed that of business-to-consumer e-commerce. This, of course, isn't news. The GartnerGroup has already reported that it expects the worldwide business-to-business e-commerce market to exceed \$7 trillion by 2004.

Call for Linux Authors

Enterprise Linux is looking for readers to write on: Linux Security, Clustering, Application Development, Linux vs. Win 2000, Network Management, Linux vs. Linux, High Availability and Disaster Recovery and Standards.

Contact Editorial Director Charlie Simpson at 215-643-8072 or csimpson@esj.com.

HP Debuts 6-Way Intel-Based Servers at 4-Way Price

In a bid to boost sales of its 4-way Intel-based servers, HP has unveiled new offerings – and challenged some conventions of the 4-way marketplace. HP's new NetServer LH 6000 and LT 6000r aren't your traditional 4-ways. They also offer 6-way processing – the capability of scaling up to six Pentium III Xeon processors.

To introduce the LH 6000 and LT 6000r, HP is giving customers the chance to order the 6-way for the same \$25,000 price as a 4-way. The promotion lasts through June. To take advantage of it, customers must purchase a 6-way design and order through HP's Select Express Program. In this program, HP's Channel Partners funnel customer orders on

to HP, which then configures, tests, and ships the machines. A 6-way NetServer carries a price tag of \$30,000. Prices for the 6000s start much lower than this, of course – \$7,299 for the NetServer LH 6000 and \$8,199 for the LT 6000r. The servers will ship this month.

In the LT 6000r, the server snaps onto a slide. Users slide the server out to upgrade and service the system.



FROM THE BENCH

HP is publicizing a number of different benchmarks that show the 6000s outperforming the Compaq ProLiant 6400 and Sun Enterprise 450. The LH 6000 equals or betters the performance of the Compaq ProLiant 6400 in a 4-way configuration and outperforms it, depending on the benchmark, 20 to 50 percent in a 6-way configuration. The LH 6000 outperforms the 450 from 30 to 80 percent in a 4-way configuration and from 60 percent to more than 200 percent in a 6-way configuration.

Along with their 6-processor capability, the NetServer LH 6000 and LT 6000r offer quite a bit of headroom. They are expandable to 8 GB of memory and offer up to 8 free PCI slots, up to 4 hot-pluggable redundant power supplies, and up to 216 GB of internal storage. They also provide two channels of embedded disk array (NetRAID) controllers at no extra cost.

The LH 6000 comes in two versions – a pedestal and a rackable version. The LT 6000r is a rack-optimized model targeted at service providers, who need to consolidate servers and externalize storage. Packaged in a 4u form factor, the LT 6000r gives ISPs and ASPs the ability to squeeze 60 Xeon processors into a 2-meter rack.

HP BEEFS UP STORAGE FOR THE MIDRANGE INTEL-BASED SERVERS

With the launch of NetServer LT 6000r systems, HP has also boosted storage capacity for the midrange by unveiling the SureStore E Disk Array FC60. The FC60 boasts high-speed Fibre Channel connectivity to NetServers and the capacity to handle heavy traffic loads and deliver fast response times for applications in areas like business intelligence, transaction processing, and e-services. The FC60 can be connected directly to servers or used in SANs. It can be reconfigured to support HP-UX.

The FC60 has a modular design. Each module contains 10 disks as well as dual controllers, for redundancy and performance. That allows users to add 10 disks at a time. The FC60 scales up to 60 disks in a 2-meter rack, providing up to 4.5 TB of main storage for the LT 6000r.

HP maintains that the price tag for this class of storage is as little as "one-tenth of what you pay for the XP256 [HP's high-end disk array]." In fact, the FC60 provides roughly 300 GB of storage for between \$50,000 and \$75,000. Contrast that with the XP256, where prices start at \$256,000.

SLIC WORKS THE MARKET

Alacritech Inc.'s new SLIC (session layer interface card) Technology has been found to improve server performance in the Internet, corporate networks and network-based applications. Alacritech will offer high-performance server adapters, delivering performance improvements in server CPU performance and network throughput.

The company's first products will be two- and four-port 10/100 server adapters, targeted specifically at enterprise users, system integrators and computer system OEMs. The multi-port design supports aggregation of multiple ports to deliver Gigabit performance, allowing customers to upgrade systems without the

inherent expense and time required to move to a Gigabit solution.

Alacritech maintains that CPUs today provide enormous amounts of power, yet often aren't able to use that speed to deliver improved network performance because they are tied up processing Internet protocols. Today's demands for greater network throughput cannot be solved using a conventional network adapter architecture, simply because they are dependant on the CPU for all of their horsepower.

Alacritech's SLIC architecture, designed to eliminate the burden of protocol processing on the server CPU, does not represent an "intelligent" or enhanced network adapter. Instead, SLIC represents a com-

plete re-thinking of the connection to the corporate LAN, via a new class of network adapter that provides higher network throughput, lower CPU utilization and greater server efficiency. SLIC Technology employs a custom Internet Protocol Processor (IPP) for protocol processing, rather than relying on the CPU to handle these tasks. Alacritech Fast Ethernet server adapters come bundled with up to 8 MB of onboard RAM for protocol processing and buffering. Conforming to PCI 2.2 specification for 64-bit devices for operation in either a 32- or 64-bit PCI environment, the products support Windows NT.

For more information, visit their Web site at www.alacritech.com. ♦

MEET THE PRINTERS HP WOULD SELL YOU IF THEY COULD.

**NEW XEROX NETWORK PRINTERS:
MORE SPEED, LESS COST AND REMOTE INTERNET PRINTING.**



PC WORLD
BEST BUY
May 2000

DocuPrint N2125

- 1200 x 1200 dpi at full rated speed
- 21 ppm vs. 17 ppm
for HP LaserJet 4050N

\$1299* or \$51/month



DocuPrint N2825

- Wide range of paper weights and sizes
- 28 ppm vs. 24 ppm
for HP LaserJet 8000N

\$1899* or \$74/month



Tektronix
COLOR PRINTERS BY
XEROX

Tektronix Phaser 850 Color Printer

- Brilliant color from solid ink technology
- 14 ppm vs. 4 ppm
for HP Color LaserJet 4500 Series

\$2495* or \$91/month

The latest additions to the new Xerox family of network printers offer you more than HP. For less. From faster speeds to effortless printing from remote locations over the Internet, you'll find everything you need built in. And every Xerox printer is backed by the Xerox Total Satisfaction Guarantee. We're so sure our printers will outperform HP, we'll give you a rebate certificate for up to \$50 for Xerox and Tektronix brand supplies if you call before June 30 to meet with your local Xerox reseller. To find out more, call 1 800 34 XEROX and meet the printers that beat HP.

Get a \$50 rebate certificate for supplies just for meeting with a Xerox reseller.

call **1 800 34 XEROX** ext 512
(349-3769)

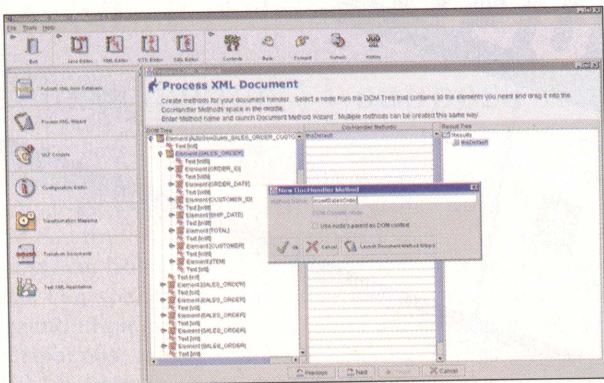
visit www.xerox.com/netprint



DIGITAL
THE DOCUMENT COMPANY
XEROX

XEROX® The Document Company®, the digital X® DocuPrint® N2125, N2825, and Phaser® are trademarks of XEROX CORPORATION. Tektronix® is a trademark licensed to XEROX CORPORATION. HP® and LaserJet® are trademarks of the Hewlett Packard Company. *Estimated street price for base model (\$US). Estimated monthly lease payments are based on a 36-month lease and may vary. Restrictions apply. Call 1 800 34 XEROX for details.

Bluestone XML for HP-UX



JAVA-BASED XML SUITE FOR HP-UX

- ▶ Bluestone Visual XML 1.1 and XML Server 6.1.5 will run on versions 10.x and higher of HP-UX.
- ▶ Bluestone XML Server 6.1.5 sells for \$2,995, VisualXML 1.1 for \$99.

Bluestone Software Inc.
1000 Briggs Road
Mount Laurel, NJ 08054
Tel: (856) 727-4600, x3008
Fax: (856) 778-8125
www.bluestone.com

It seems that 1999 was the year for extensible markup language, or XML. Beginning with Microsoft's Q1 announcement of an e-commerce strategy that heavily leveraged an XML framework, the little-language-that-could just plum took off. With its own late-year announcement of a Java-based XML Suite, Bluestone Software Inc. hopes to bring a comprehensive XML solution to a variety of platforms – including HP-UX.

XML is often erroneously equated with Java, a programming language and runtime environment that boasts the ability to compile and run on just about any platform for which a Java virtual machine has also been written. Rather than simply providing the capability to run on heterogeneous platforms, XML's real value-add is its ability to "glue" disparate platforms together. And XML's "glue" doesn't simply extend to communications

between platforms themselves, but, rather, to the business applications that reside on these platforms.

For its part, Bluestone's XML suite is composed of version 1.1 of Bluestone VisualXML, an XML development environment and toolkit for building XML applications. Visual XML 1.1 is matched on the server side with version 6.1.5 of the Bluestone XML Server. According to Al Smith, Vice President of Software Development with Bluestone Software, both Visual XML 1.1 and Bluestone XML Server work to complement one another.

"VisualXML is a tool that has two purposes, one of which is to let people build applications that would be hosted in Bluestone XML Server and execute on a runtime basis, and the other aspect is a client-side tool that actually lets you manipulate XML," says Smith. "And Bluestone XML Server is, of course, the platform on which you can host your XML applications, as well as interoperate with other XML server hosts."

Because of its inherent intuitiveness and self-documenting facility, Smith acknowledges that XML will likely play a large part in the burgeoning e-commerce space. But there's another potential use for XML that Smith says could be just as big – as a replacement for expensive, proprietary and often unwieldy electronic data interchange (EDI) solutions.

"XML lets you be unconstrained by somebody's previous definition," Smith says. "So when you look at EDI, a lot of times when I get done mapping [everything] from my AS/400 ... to my HP 3000 or 9000 system, I'm left with some data that just won't come over," he explains. "But when you look at XML, a lot of that excess time goes away because there's a definition of the data with a data, and so it's very easy to modify your programs."

And as enterprises increasingly gear up for e-commerce or seek to implement next-generation EDI solutions, Smith says that a back-end platform like HP-UX can benefit from XML's ability to "glue" it to other operating systems and application platforms, across the divisions of a company or in business-to-business scenarios.

Smith's optimism is echoed by members of the analyst community as well. Tim Minahan, a Senior Analyst in electronic commerce with consultancy at the Aberdeen Group, says that XML-based technologies – like Bluestone Visual XML and Bluestone Server – can help to tie together heterogeneous platforms.

"Aberdeen believes that XML technologies have the ability to provide support for open application integration, allowing disparate systems to interoperate with one another," he affirms.

– Stephen Swoyer,
Contributing Author



SEE THE BIG PICTURE...

YOUR LAN, YOUR WAN, THE WORLD

IBM Sun Microsystems Microsoft NT hp HEWLETT PACKARD digital Novell

Browse And See The Status Of All Your Servers... From Anywhere In The World.

NIGHTWATCH NT™

The "Central Console" Piece Of The Nightwatch Enterprise Solution.

Nightwatch NT is a failsafe event notification system. It scans a list of specific software and hardware events you tell it to monitor.

Nightwatch alerts you when triggered by problems you want to know about such as...difficulty with system backups, disc space concerns, WEB PAGE availability, ERROR STRINGS found in PROCESS LOG FILES files, and much, much more!

Nightwatch NT can 'listen' for problems or 'go get' information from all different types of application software, and all different brands of hardware and operating systems that run throughout your Enterprise.

NIGHTWATCH®
INSTALLS IN 5 MINUTES!

**You Can Be Monitoring
Important Events In
Your Enterprise In Less
Than 30 Minutes!
CALL AND WE'LL SHOW
YOU HOW!**

Nightwatch supports alphanumeric and numeric paging, as well as cell phone, E-MAIL and WINPOP notification.

Hillary Software, Inc.

309 Morris Avenue, Spring Lake, NJ 07762 USA e-mail:sales@hillary.com
1-800-HILLARY 1-800-445-5279 1-732-974-8484 FAX 1-732-974-0167


Visit our web site www.hillary.com for up-to-date technical details!

Nightwatch & Nightwatch NT are products of Hillary Software, Inc. All Rights Reserved ©1998-2000Hillary Software, Inc.
Other brands and/or company names that appear herein are registered trademarks of their respective companies. They appear for identification purposes only and do not imply any endorsement by them.

Smart Ways of Managing Data

Data Warehousing and Business Intelligence

Jean Nattkemper



No one's complaining about a shortage of data. In fact, data is accumulating so rapidly, and in so many different sources, that organizations are struggling to consolidate the information and provide access to it across the enterprise and within the supply chain. One way of dealing with data consolidation and access is, of course, by creating a data warehouse (DW). DWs offer a means of collecting, cleansing, transforming and storing transactional and external information in a structured format for query analysis and reporting. But, they are just one tool for managing data.

Thanks largely to the Internet, the nature of data has changed. Today, about 90 percent of data exists in unstructured formats – in word processing programs, e-mail and so on. That's one reason why organizations are looking beyond data warehouses to business intelligence (BI) software and solutions. BI allows organizations to address specific business questions by accessing and analyzing information from a DW, or from other structured or unstructured data sources.

HP 3000, HP 9000

Users Get Reports from LANS, WANS, the INTERNET or as EMAIL Attachments.


Reports Arrive In Formats Like These:



WORD EXCEL HTML PDF

- **byREQUEST**™ Gives Your Users Just What They Want, When They Want It, In Useful Formats They Understand!
- Distribute Reports In 'Real' Time, Locally Or Worldwide.
- View And Archive From CD.
- View Reports In Word, Excel Or PDF.

byREQUEST
INSTALLS IN 5 MINUTES!
You Can Be Formatting And
Delivering Your Own Reports
In Less Than 20 Minutes.
Download A Demo, Call For
A CD Key... And We'll
Show You How!



epaper
FROM LEGACY
MACHINES™

byREQUEST™

Does It
Automatically!

DISTRIBUTE REPORTS USING PC FORMATS USERS UNDERSTAND

SPREAD THE WORD...IN WORD...IN EXCEL OR PDF.


byREQUEST™

Hillary Software, Inc.

309 Morris Avenue • Spring Lake, NJ 07762 • USA

1-800-445-5279 • 1-732-974-8484 • Fax: 1-732-974-0167

Visit us on the web at www.hillary.com for demonstrations

Address  <http://www.hillary.com/>



epaper
FOR BUSINESS™

byRequest is a product of Hillary Software, Inc. All Rights Reserved ©2000 Hillary Software, Inc.

Other brands, icons and/or company names that appear herein are registered trademarks of their respective companies. They appear for identification purposes only and do not imply any endorsement by them.

In any case, BI/DW presents huge opportunities – Survey.com reported that BI/DW became a \$25-\$30 billion market worldwide last year. While IBM remains the dominant player here, offering a complete mix of hardware, software and services, HP seems poised to capture substantial market share.

One of HP's big advantages is hardware. UNIX is still the number one platform for data warehousing, according to IDC, and HP-UX servers provide the scalability needed to run BI/DW solutions.

HP is also taking advantage of its partnering expertise to forge new alliances in the growing enterprise information portal (EIP) space. Portals are essential to BI because they bring intelligence to "the so-called masses in an organization," says Thomas Wuerz, HP's Business Intelligence Program Manager. Because of partnerships it has forged with vendors, like Cognos, Brio and Viador, HP can offer EIP solutions that allow information previously available only to a few specialized employees to be distributed and accessed throughout an organization.

In perhaps its most interesting – and certainly its most innovative – initiatives, HP is developing BI solutions predicated on e-services. Portals also play a central role here, as does e-speak, HP's brokering technology for dynamically creating e-services. HP thinks it's on to a good thing with e-intelligence. "Our e-intelligence initiative isn't just business intelligence with a Web interface slapped on it," says Jennie Grimes, HP's Director of e-Intelligence and Portal Solutions.

As HP moves forward on different fronts in the BI/DW space, partnering remains a core strategy across initiatives and programs. Take, for example, the HP Open Warehouse Program, which was launched in 1993, and by 1998, had evolved into the HP Business Intelligence Program. This partnership program includes major initiatives in the areas of ERP/CRM, high-end data warehouses and EIPs.

HP, SAP'S ERP/BI INITIATIVE

"When HP talks about ERP, there's a strong focus on SAP and the SAP Business Information Warehouse (SAP BW)," Wuerz says. SAP BW is the ERP vendor's initiative for BI solutions in an R/3 environment, as well as the underlying piece for new SAP initiatives, like mySAP.com (its portal solution), CRM applications, and procurement solutions.

BI solutions for a SAP R/3 environment are a natural for HP, which boasts more than 1,000 installed SAP BW solutions and more than 6,000 installations running R/3. "HP's value proposition in this market," Wuerz explains, "is a combination of hardware, storage, tools and services to help businesses obtain vital reporting and analysis out of their SAP data."

HP delivered on that business proposition by unveiling the Impact Analysis for SAP BW. The analysis shows businesses how SAP BW can complement a current or planned data warehouse architecture, and identifies and costs all the resources needed to implement and operate SAP BW successfully. HP also developed a configuring and sizing tool, called CAST, that, Wuerz says, is used to "effectively reduce the risk involved in sizing a SAP BW solution, while making it easy to estimate the hardware needed."

HP also offers the Certified Data Warehouse program for SAP. Forged out of an alliance between HP, Informix and Cognos, the program is designed to reduce the time it takes to build an SAP BW solution. Program partners test and certify SAP BW solutions for high-performance warehouses in three configurations – small, medium and large.



HP's BI Program Manager, Thomas Wuerz

company will use the warehouse mainly to analyze call detail records (CDRs) for the purpose of determining call behavior trends.

The Cap Gemini deal isn't an anomaly in the market. More and more businesses consider data warehouses in the multi-terabyte range a mission-critical necessity. In fact, IDC predicts that by 2001, 5 percent of the data warehouses installed will hold more than 16 TB of data, and 20 percent will hold more than 2 TB.

Again, HP is relying on partnerships to serve this market. An essential partner is Oracle. Toward the end of 1998, HP and Oracle launched the Terabyte Plus initiative to build warehouses that can be accessed by hundreds to thousands of enterprise and Internet users. The two companies' combined offerings are designed to "look as though they come from only one company," Wuerz says.

Under terms of the partnership, HP and Oracle integrated HP 9000 data warehouse servers with the Oracle8i database and established an Enterprise Technology Center in Atlanta to help customers design and deploy multi-terabyte data warehouses. The partnership was also marked by an industry "first" when the companies aligned their separate support organizations in the Service Integration Program.

HP, ORACLE TERABYTE PLUS

HP knows all about the need for large data warehouses. In one huge deal, HP agreed to build a 40-TB data warehouse for Cap Gemini over two years. The telecommunications

HP ENTERS EIP MARKET

HP's interest in the EIP market is fueled by the two big advantages of these portals (see Figure 1). First, EIPs allow for wider distribution of information within an




Figure 1: Web-Based Access Through EIPs – EIPs provide a single point of Web-based access to enterprise information while searching across both structured and unstructured sources of data.

CAN YOU FIND THE DATA IN THIS PICTURE?

IT'S EASY. THERE ARE ALMOST 200 TERABYTES OF DATA ON THE IBM MAGSTAR® VIRTUAL TAPE SERVER, AND LIKE MAGIC, YOU CAN BRIDGE THE DATA TO A SECOND SYSTEM IN A REMOTE LOCATION. WITH ITS BUILT-IN INTELLIGENCE, THE IBM MAGSTAR VIRTUAL TAPE SERVER PROVIDES LOGICAL VOLUME STACKING FOR NEAR-100% TAPE CARTRIDGE UTILIZATION — RESULTING IN REDUCED TAPE OPERATION COSTS AND MORE TAPE DATA UNDER AUTOMATED CONTROL.

LEARN MORE ABOUT THE FUTURE-READY IBM MAGSTAR VIRTUAL TAPE SERVER AT ibm.com/storage/magic/vts

The IBM logo is displayed in its classic eight horizontal stripes. It is positioned on the left side of the advertisement, overlaid on a background image of a city scene with a bridge and water. The logo is white with a black outline.

IBM @ business storage servers.
Technology. Innovation. Magic.

IBM, Magstar Virtual Tape Server and the e-business logo are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries. © 2000 IBM Corporation. All rights reserved.

organization. And second, their use of search engines allows them to offer a means of analyzing unstructured data – the data residing outside databases, data warehouses, ERP systems and data marts.

“All of the BI vendors are moving into the EIP world to add collaboration or the capability of analyzing unstructured data to their offerings,” says Wuerz, who points to Cognos, Information Advantage (the OLAP software vendor acquired by Sterling Software last year), and Viador as examples.

Viador, in fact, is a major HP partner in the BI space. “About 80 percent of Viador’s business today is aimed at employees in companies who are analyzing data sources that exist within the company,” says Wuerz. “That’s very different from other portal trends, but something all other EIP vendors are morphing into.”

Viador’s E-Portal Suite, which has embedded within it Infoseek’s Ultraseek Server enterprise search application, allows an organization’s users to access and analyze information across databases and documents. Last October, HP and Viador forged a strategic alliance to deliver an online implementation service that enables the rapid deployment of an EIP.

INTELLIGENCE MORPHS INTO E-SERVICES

HP took great strides to move out from its more traditional offerings in the BI/DW space last December when it launched its e-intelligence initiative. The company began talking about e-intelligence, says Jennie Grimes, HP’s Director of e-Intelligence and Portal Solutions, “When we tried to see what the Internet was going to do to the traditional business intelligence world. And we figured out that there wouldn’t be just an evolution. A revolutionary set of activities would emerge.” The Web interface stage is just the beginning of entry into the e-intelligence world, Grimes says, although she expects brick-and-mortar, or chapter one Internet customers to ease themselves into the e-services world by taking what is “almost a first-generation step.”

“When we look at the end state of e-intelligence, what we see are a couple of key things that tell me it’s a functional shift than rather a little step-wise movement,” Grimes explains. “What we realized first was that in this e-services world, as our customers try to figure out what to sell, how to sell it, to whom to sell it, for how long to sell it, and at what

Applying BI to Internet Usage Data

HP entered the emerging IP usage metering market a couple of years ago with a product called Smart Internet Usage (SIU). SIU is a layer of software that sits between the IP network and business support system of a service provider, or an enterprise. It collects information on resource usage from different network elements and packages that data into Internet Data Records (IDRs) for use by billing systems, data mining applications, or operations management applications.

The earliest targeted application for the kind of information provided in IDRs is usage-based billing for ISPs. But HP’s Dana Kreitter, Marketing Manager for SIU, sees “an equal, if not greater, opportunity for SIU on the BI side.” HP is positioning SIU to service providers as technology that enables them to gather the kind of information that will allow them to offer differentiated services and reduce customer churn. Using data gathered by SIU, service providers can segment their subscriber base and create service bundles based on customer interests. Differentiated service offerings allow providers to step out from the pack of ISPs that offer one-size-fits-all approaches. They also give providers a way of enhancing customer satisfaction to reduce churn – loss of existing subscribers.

Because SIU is just the middleware layer that collects usage information and makes it available, HP is forging partnerships with vendors that provide modeling solutions for the churn problem or that offer data mining applications that contain ad hoc query interfaces. Partnerships open up “broad possibilities,” Kreitter says. “For example, we might produce our data records and stuff them into somebody’s big Oracle database. Then they can use their business modeling to combine the IP usage information with other information, perhaps demographic data or data about other telecommunications services.”

SIU offers some other possibilities, according to Kreitter. Service providers, for example, often determine how to price a new service by operating it on a small-scale trial basis. SIU can help provide information on uptake – the customers who are making heavy use of the service – while also providing usage information from a traffic engineering standpoint to determine the cost of service delivery – the percentage of bandwidth that will be used by this new service.

Kreitter acknowledges that IP usage metering is an emerging market not only for HP but for the industry in general. “Only in the past year has there been a significant uptick in interest [in these types of applications] from the heavy hitters – the large-scale service providers.” But Kreitter is in the process of lining up application partners to fill in the gaps so that HP can offer SIU as part of an entire package enabling data analysis.

– J.N.

price to sell it – all of that is fueled by intelligence. All of it is fueled by my supply chain information, and that’s fueled by my customer relationship information, and that’s fueled by my current business performance information. And, we said, ‘Look, that’s intelligence seeding e-services.’”

HP realizes that’s not how BI is viewed now. To most businesses, BI is seen primarily as a downstream reporting activity. In an e-services world, though, everything is dynamic. And BI is an inherent part of every transaction.

“Look at it this way,” Grimes says. “Every transaction has intelligence embedded in it, and every transaction had intel-

ligence that drove it.” Companies already sell intelligence – psychographic and demographic profiling, for example. “But,” Grimes says, “what if I’m selling the clickstream analysis or what if I’m selling the customer profiling in the pharmaceutical segment? What if I’m selling the data extract capability? What if I’m selling the data mining capability or the cleansing or the scrubbing – all the parts of a traditional data warehouse?” At this point, intelligence becomes an e-service.

E-INTELLIGENCE THROUGH ISV PARTNERSHIPS

It’s one thing, of course, to talk about e-intelligence, and quite another to deliver

on it. But, HP is pulling together a lot of the pieces that make the initiative seem real and compelling.

One of these pieces involves ISVs, who play a major role in HP's e-intelligence efforts. Last December, HP launched the E-Intelligence Partner Program, which initially embraced more than three dozen ISV partners, including Acta, Hyperion, Infoseek, Informatica, e.phiphany, PeopleSoft and SAS. The new partner program welcomes traditional, or brick-and-mortar, ISVs, like SAS and Oracle, which have business models built on software license transfer.

For example, HP announced that it had collaborated with SAS on an eCRM knowledge portal, leveraging SAS' decision-support technologies. The knowledge portal combines SAS data mining solutions and services with HP hardware, support and consulting to educate customers on conducting customer evaluation, campaign analysis, and segmentation profiling.

Dotcom ISVs, companies just getting started in the BI space, are also embraced by the E-Intelligence Partner Program. HP has, for example, developed an e-service with ShortCycles. The application is designed to replace intranet and extranet sites with a single solution for exchanging realtime information between sales and marketing departments, business partners, and key customers.

The program is also platform-independent insofar as HP platforms go. In other words, ISVs receive assistance in porting applications not only to HP-UX, but also to MPE/iX, Windows NT and Linux.

HP's ISV partners are expected to embrace e-speak as part of their strategic direction. Brick-and-mortar ISVs must work within product upgrades and release cycles to find the appropriate release in which to embed e-speak. Dotcom ISVs, because they're just getting started, can snap e-speak into their architecture rather than retrofitting it.

HP e-intelligence partners that are piloting e-speak projects include Ericsson, VHx, Hitel and Jobpartners.com. Ericsson is using e-speak in the area of service repair dispatch on a solution whereby a request for services will be routed to the appropriate technician via cell phone. Healthcare company VHx, just acquired by Claimsnet.com, is using e-speak to enable its healthcare exchange. Korea-based Hitel is



HP's Director of e-Intelligence and Portal Solutions, Jennie Grimes

using the technology in its portal solutions for online gaming. Finally, Jobpartners.com is using it to develop a recruiting portal.

HP also is enticing ISVs by providing them with a portal through which they can tap into HP's resources. The portal, which provides information on customer successes, additional solution sets and additional feature functionality, is an avenue

open to HP customers, system integrators, sales reps and "really anyone," Grimes says.

PULLING THE PIECES TOGETHER

E-intelligence reaches far beyond ISV partnerships to embrace new service delivery models (like Viador), financing options, support, storage, ASPs and new intellectual property yet to be developed.

"When looking at financing options with ISVs and, in some cases, with customers, we ask ourselves, for example, if the relationship is with an ASP? Is this a lease space? A click space? How do we fund it? And we get HP Technology Finance involved," says Grimes.

As far as support models are concerned, HP is finding that e-intelligence creates its own support needs. Grimes and her group, for example, are working with HP's support organization to create a program similar to HP's recently announced critical systems support program for e-commerce. The new program will map specifically to e-intelligence.

E-intelligence also affects storage and storage support. A business running an ERP e-commerce application may be able to cache off last week's transactions every day, says Grimes. In that case, 500 GB is a lot of storage. In the e-intelligence world, though, businesses engage in activities like clickstream analysis, causing storage requirements soar. "Each week Amazon.com generates 500 GB of clickstream data alone, and that's not counting all the repositories they're caching off to do customer profiling and trend analysis," Grimes says.

"An HP support contract in a traditional area would concentrate very heavily on the high availability of the server, for example, and MCServ/Guard failover strategies," Grimes points out. "We'd offer a support contract that guarantees six-hour repair on servers." An e-intelligence application is an entirely different matter. "It's actually the storage

you're worried about," Grimes explains. "You can failover your servers, but if you don't have a good robust disk array or three-way mirroring on the storage side, and you take an Oracle database corruption hit, it can take you 14 to 18 hours to cycle through the corrupted database." In other words, support contracts must be adjusted to focus not only on the servers, but on the kind of storage area network (SAN) strategy used to help keep the storage devices in synch and available.

EXPLORING UNCHARTED TERRITORY

As HP moves into uncharted territory, driving intelligence as an e-service, it is focusing on a couple of specific areas. One is the notion of a single customer view, or the integration of data from all customer touchpoints. "HP has flirted with this notion for years through key programs like our CRM program and our CBA, or Call Behavior Analysis, program," Grimes says. "One [program] stemmed largely from the financial community, and the other from the telecommunications world, but they were both after that single customer view. Now take something we're working on now – the virtual single customer view." With the virtual single customer view, businesses may use business process technology and create a ripple effect in which data is distributed not only from a CRM application, but also from a supply chain application, or perhaps some of the operational systems.

HP is also expanding its portal activity for e-intelligence. Along with the EIP arena, the company is pursuing partnerships and solutions for e-commerce and trading community portals.

The company sees yet another opportunity in the BI space. It hopes not only to leverage existing intellectual property, like e-speak, but to foster development of intellectual property that will enable e-intelligence applications. "This is a real opportunity for HP," Grimes says. "We think we're out ahead of the industry here and able to develop product – actual software product – that we think can uniquely drive some e-intelligence solutions." ♦

— Jean Nattkemper is the Editor at Large for HP Professional.

She can be reached at jnattkemper@hpro.com.



How to Mine Your Own Business

Mining for Gold in a World of Changing Currency

Katherine Hammer

The conviction, "There's gold in them thar hills!" is at the heart of both data warehousing and data mining.

The difference between the two is in how this "gold" is obtained. Data warehouses are valuable when managers know what variables should be tracked. While this information can be extremely valuable in tuning the way a business is run, too often the results and their interpretation are colored by the way the questions are asked.

The goal of data mining is to look for some correlation between the variables stored in a warehouse that is not expected, some aberration, which upon further examination could lead to some new insight. Data mining technology takes information about how the elements in a warehouse are related and uses technology grounded in statistics and neural networks to look for patterns between values that could be significant.

THE PROBLEMS OF DATA MINING

Data mining technology has several important prerequisites, namely:

- The data in the warehouse is clean and consistent
- The individuals deploying the data mining products have correctly specified the relationships between variables
- Management defines criteria for what it will spend to achieve "explanatory adequacy"
- The warehouse architecture supports evolution to facilitate the acquisition of


any additional data

Clean and consistent data.

If the implementation team cannot ensure that the data represented in a warehouse is consistent, mining tools may turn up interesting correlations between data elements that are totally spurious. The whole value of the data mining application is compromised.

The complexity of the task of populating a warehouse with clean and consistent data is a function of how badly the warehouse is needed. If the goal is to capture a subset of data from an operational database and transform the data values so they can be used by business users in a data mart, the task of ensuring clean and consistent data is fairly straightforward. The implementation team merely needs to ensure that the source database is consistent. However, even when dealing with a single database, this task is not simple.

Frequently, it is not possible to rebuild historical data in a database used by a mission-critical system because there is not sufficient downtime. As a result, particularly in legacy environments, a DBA will make a schema change in such a way as not to alert the data manager. Sometimes these changes are affected by COBOL REDEFINES, sometimes by a particular data value that signals the existence of another file or record. Because these changes are rarely documented in one place – rather, they are distributed across file descriptions, DBMS data defini-



tion files and application code – extensive data sampling may be required.

However, when a company is building a warehouse that either consolidates equivalent data from multiple sources (as in the case of pulling data from multiple purchasing systems) or merges related data from different types of applications, the task of populating the warehouse is even more complex. Not only must one deal with inconsistencies within a particular source database, but one must recognize when the same customer or vendor is being referenced across multiple source databases, create new keys, etc.

Despite these complexities, if a company is not committed to making the investment required to ensure clean, consistent data, it might as well give up any thought of using data mining technology.

Accurate identification of dependencies. To reduce computational complexity, data mining products allow the user to specify the relationships between the elements in a warehouse. In many respects, these are not unlike the kind of relationships captured in ERA diagrams used in designing databases. However, when defining a database, one is defining such things as the legal relationship between entities and

▼ Knowledge-Driven

**Schooled in the
art of knowledge
transfer.**

- ▶ When it comes to business intelligence consulting, Enterprise Group, Ltd. takes a unique approach. We offer the knowledge and teach the skills you need to manage your business intelligence systems. Rather than encourage dependence, we foster self-sufficiency. In other words, rather than cook you a fish dinner, we teach you to fish.

Enterprise Group, Ltd.'s client-centered approach facilitates knowledge transfer. We provide side-by-side mentoring and help you cultivate strong cross-functional project teams. We host structured classes and develop extensive on-site libraries. And through our affiliation with the BI Alliance, we tap the brightest minds in the industry as your resource.

We help you harvest the deep waters of information management and focus on long-term success.





Enterprise Group Ltd. helps organizations maintain a global competitive edge by providing a full range of business intelligence and information management services.



**ENTERPRISE
GROUP, LTD.**

We build business intelligenceSM



www.eg ltd.com 800.428.2005

the legal data types or range of values for each attribute.

The DBMS guarantees that these specifications will be enforced. In contrast, defining parameters for use with data mining technology is more a matter of heuristics, that is, a set of guidelines for what is expected. Data mining products use this information to "notice" relationships that violate these heuristics.

One value may be dependent on another – for example, line of credit may be a function of both credit rating and income. A higher line of credit implies that payments will be timely. A data mining tool might flag a potentially important correlation if it finds a population of customers with a pattern of high lines of credit and a large number of late payments. What this correlation might mean is a function of how this correlation is interpreted. Depending upon the company processes used to derive the values in the warehouse, there could be very different interpretations. If line of credit is computed automatically and enforced by some computer program, then it is unlikely that there is an error in line of credit. If, on the other hand, employees are the ones who make this determination and enter the line of credit, management might want to see if these customers are tied to some geographic location serviced by a particular outlet of the company, and so on.

In short, the confidence with which management can be certain that it has found a real nugget from data mining is no easy matter. It is a function of the accuracy with which these heuristics have been specified; management's knowledge of the business processes and systems in place; the ability to add additional data to the warehouse in case the correlation is important enough that it should not be ignored and a number of competing interpretations are possible.

Determining explanatory adequacy. Scientific inquiry is a good analogy for the dilemma of companies trying to benefit from data mining. In trying to understand a phenomenon, it is as important – and perhaps easier – to discount the variables that don't contribute as to understand the ones that do. As Lou Agosta writes about data mining in *The Essential Guide to Data Warehousing*, "Refutation is absolute; whereas confirmation is always partial and tentative." Just as the early stages of scientific investigation typically lead to more questions that lead to more experiments, data mining technology is a tool for helping management refine the questions it

should ask. Sometimes sufficient evidence can be found from these queries to point to a probable conclusion.

The dilemma comes when all obvious questions have been posed to the warehouse, and no definitive interpretation has been found for the correlation uncovered by the data mining. There are two possibilities at this point: Conclude that the

correlation is irrelevant or generate a series of hypotheses, or theories, about what could account for it.

If the operational systems in the organization contain data that could provide corroboration for one or more of the competing hypotheses, management and IT must assess how quickly information

See Mining on page 33

Eleven Steps to Success in Data Warehousing

Phillip
Blackwood

Although data warehousing has long been an option for big companies, the reduction in warehousing technology costs now makes it practical for smaller companies as well. Turnkey integrated analytical solutions are reducing the cost, time and risk involved in implementation, and corporate portals are making it possible to grant access to hundreds, or thousands, of employees.

1. *Recognize that the job is harder than expected.* Experts report that 30 percent to 50 percent of the information in a typical database is missing or incorrect. In an operational system that focuses on swiftly and accurately processing current transactions, this may not be noticeable or may even be acceptable. However, it's totally unacceptable in a data warehousing system designed to sort through millions of records in order to identify trends or select customers for a new product. Another challenge is that although database schema changes over the life cycle of a project, historical databases are rarely rebuilt.

2. *Understand the data in existing systems.* It is important to perform a detailed analysis of the status of all databases that will contribute to the data warehouse. Data interrelationships need to be determined between various systems, and must be maintained as the data is moved into the warehouse. Since the data warehouse implementation often involves making changes to database schema, understanding data relationships among heterogeneous systems is required to predict the impact of any such change, and avoid inconsistencies in other areas of the enterprise.

3. *Recognize equivalent entities.* One of the most important aspects of preparing for a data warehousing project is identifying equivalent entities and heterogeneous systems. For example, two different divisions may be servicing the same customer, yet have the name entered differently in their records (i.e., "AIG" and "American International Group"). A data transformation product capable of fuzzy matching can be used to identify and correct this and similar problems.

4. *Support data quality with metadata.* Metadata is data about data; for example, tags that indicate the subject of a World Wide Web document. One major challenge is trying to synchronize the metadata between different vendor products, different functions and different metadata stores. It is important to create and capture metadata for interfaces, business processes and database requirements as soon as possible. Several vendors offer products that have the potential to integrate metadata from disparate sources and begin to establish a central repository that can be used to provide the information needed by both administrators and users.

5. *Select the right tools.* Data transformation tools extract data, clean it and load it into the data warehouse, while capturing the history of that process. This transformation process may include the creation and population of new fields from the operational data, summarizing data to an appropriate level for analysis, performing error checks to validate data integrity, etc. Look for a tool that can map data from source to target with a point-and-click interface. The ability to track and manage the relationships of interrelated data entities is also useful. Finally, try to find a tool that can capture and store metadata during the conversion process.

See Eleven Steps on page 33

Data Warehouse Factors to Address for Success

META Group's Scorecard Report Illuminates Key Factors

Doug Laney

The responses and subsequent analysis in META Group's 1999 Data Warehouse Scorecard Report confirm that team experience, architecture, expectation management and data quality are the prevailing factors influencing the success of data warehouse (DW) initiatives.

More than half of respondents report moderate success or better, although outright, documented return on investment shows up for only one-third of respondents. With the emerging knowledge base of DW best practices and a growing skill base of DW practitioners, fewer than 10 percent of DW installations are now utter disasters.

The availability of hard "in-practice" data, such as this report, gives DW project sponsors, managers and architects a tool to benchmark leading technologies and approaches against DW success rates. With these DW characteristic reference points, existing project teams can rationalize their levels of success, and newly formed project teams can improve their odds of success. By 2000/01 these success factors will be firmly ingrained in the DW community's consciousness, and by 2003/04 the rate of reported DW installation success will improve to con-



verge with rates for traditional application development initiatives.

Through 2003/04, essential DW success factors will remain unchanged even as new architectures, strategies and technologies emerge. Sets of best practices and DW standards will form around these success factors, leading to the formalization of DW projects well beyond current experimental tendencies. Built-in DW capabilities currently found in some packaged business solutions will be evi-

dent in all leading packages by 2001/02 and will reflect these winning architectural principles to a large extent. Notable non-factors in determining DW success include the organization's industry type and the choice of outside systems integrators.

The traditional DW concept has evolved into a broad set of architectural variations. Data marts (standalone and federated), operational data stores, exploration warehouses and application data stores, in addition to conventional hub-and-spoke DW, have enabled sophisticated information management solutions to meet exacting business requirements. Along with leveraging the volume of data generated by ERP applications, there is a burgeon-

ing need for the DW to sustain customer relationship management (CRM), supply chain management (SCM) and electronic commerce initiatives.

META Group's 1999 Data Warehouse Scorecard Report is a point-in-time survey of existing data warehouse and data mart environments. It focuses on what is currently working in practice, and addresses the aspects of a maturing technology from the subjective perspectives of those deploying it.

ARCHITECTURE

The distribution of DW architectures between data mart-oriented (34 percent) and centralized DW-oriented (34 percent) approaches (including hub-and-spoke) is especially uniform. Though there is little preference for one approach over another, the profiles of organizations adopting them are significant. Organizations reporting failed DW implementation are two or three times more likely to have employed data mart-oriented approaches, while classic DW architectures revolving around a centralized data store enjoy a strong lead among acknowledged successful implementations.

We also find that, because single-subject data marts fail more often than formal multi-subject DWs, an insulated data mart project is no guarantee of success. For central DW environments, incremental cost of each additional subject area is constant. In data mart-oriented environments, however, incremental cost of each additional mart increases. In-production DWs employing data mart-oriented architectures also show a 70-percent higher average cost per subject area than those with centralized architectures. Our qualitative analysis suggests this stems from the exponential administrative overhead of decoupled data marts and the technical burden of ensuring consistency and synchronization of federated data mart environments.

In addition, the low rate of operational data store architectures (12 percent) indicates relative disinterest in tactically oriented information solutions. The higher rate of mixed DW architectures (20 percent) illustrates a level of information supply chain maturity for successful DW projects and confusion/politics for unsuccessful ones.

EXPERIENCE

Unsurprisingly, the longer an organization toils at DW, the more likely success will result (e.g., after 2+ years, projects are likely to have twice the level of success as 6-month-old projects). The earliest adopters tend to rate significantly higher, representing not only brute-force success over time, but also cultural propensity toward innovation and acceptance. Successful shorter-term efforts are characterized by a high degree of reliance on technology and larger team sizes. The high number of failed efforts that have persisted for three years or longer begs deeper quantitative consideration, but

we believe these results stem from politically mired projects, over-scoped (non-iterative) efforts, sponsor turnover and DW "re-dos."

Although it is not evident that larger teams are more successful, we find that undersized teams can really hurt a project. Core team sizes of seven to 10 FTEs seem optimal to avoid failure. We also find that management's common belief that a small data mart implies a small team does not hold true.

A lack of expectation management remains a barrier to DW success. Even successful DW initiatives seem haunted by the ghosts of IT projects past. However, successful efforts show double the propen-

Core teams of seven to 10 FTEs seem optimal to avoid failure.

sity to deal directly and continuously with issues like user participation, executive involvement, and overall project promotion. Within IT ranks, positive expectation management includes education on architectural, managerial and usage distinctions of DW versus traditional OLTP systems development projects.

Successful DW efforts are nearly three times more likely to deal effectively with data quality issues compared to failed projects. Organizations that jump with both feet into DW data modeling efforts neglect many critical analysis activities that deal with assessing data source availability, accessibility, suitability, "integratability," completeness and accuracy.

OTHER FACTORS

Industry. An organization's industry type does not drive DW success, though varying adoption rates by industries yielded a range of overall scoring results. For example, retail and education organizations (overall DW laggards) are experiencing slightly lower rates of success, while transportation and the historical technology leader, telecommunications, achieve success more often. Overall, however, success rates are relatively consistent across industries.

Consulting. The study also unearths the realization that DW boutiques and DW vendor consultants are found twice as often on successful DW initiatives than

large systems integrators. This finding is as much a reflection of global integrators' current inability (or lack of desire) to capture, package and share DW best-practices as it is a reflection of the prevalence of targeted competencies throughout the smaller and boutique firms. Until global integrators master the economies of scale they have enjoyed in OLTP systems development, we believe individuals, not firms, will continue to make the difference in DW.

ETL Tools. While extract-transform-load tools did not correlate highly to success or failure, highly successful organizations did tend to persist with these tools rather than giving up on them. This finding illustrates the long-term benefit of ETL tools despite their lingering learning curves.

DW project leaders and LOB sponsors must look closely and objectively at the characteristics of their initiatives. While there is no formula for DW success, attention to fundamental factors can lead teams away from danger and toward success. Organizations should apply these survey results to select partners, coordinate activities and establish DW standards. Architectural success drivers illuminated by this study validate the hub-and-spoke DW approach and data mart-centric promoters. Indeed, although technology has advanced to support various data mart-oriented approaches, these approaches are proving less effective. Pressures for direct user access and extreme scalability are borne out in the survey results.

Further, IT/LOB conflicts are being exacerbated by uncoordinated approaches, poor expectation management and underestimated data quality challenges.

Early DW adopters were encouraged to experiment broadly with various architectures, technologies, team compositions and LOB involvement schemes. But project sponsors now require that DWs be treated as a mission-critical information supply chain component. To mitigate risks of DW failure and elevate chances for success, project teams must be guided by aggregate results and insights from previous DW projects. ♦

— Doug Laney is META Group's Senior Program Director for its Application Delivery Strategies Service, leading META Group's data warehouse-related research agenda. He also serves as DCI's Business Intelligence World conference chairman.

A One-on-One with Bill Inmon

Charlie Simpson

Bill Inmon is the Chief Technical Officer of Pine Cone Systems. With over a decade of DW experience and 39 books published and translated into nine languages, Bill has been dubbed the "Father of the Data Warehouse."

HP Professional: How has today's data warehouse changed from a few years ago, in both its use and design?

Bill Inmon: Today's data warehouse is significantly expanded from the warehouse of a decade ago. A decade ago, data warehousing was a separate database, apart from the transaction database.

Today's data warehouse has blossomed into a full-blown architecture called the "corporate information factory." The corporate information factory has its own dynamics and its own component parts. There is the enterprise data warehouse, which contains the granular data, that is the foundation on which everything else depends. From the enterprise data warehouse emanates the data marts, the exploration warehouse, alternative storage – near line storage and secondary storage, project warehouses. Then, there is the operational data store (ODS) and the integration and transformation layer. In short, today we have a full-fledged framework where yesterday we had a raw, detailed database full of integrated, historical data.

HP Pro: How has business intelligence (BI) impacted the data warehouse?

Inmon: Business intelligence has enabled the end user to take advantage of the data warehouse. Data warehouse and BI go together, hand in hand. The data warehouse is the infrastructure and the BI com-

ponent is the end user access and analysis component. They form a symbiotic relationship.


HP Pro: Is outsourcing data warehousing development and maintenance a real option for IT managers?

Inmon: The appeal of outsourcing the data warehouse development is great. Data warehouse skills are very difficult to find and expensive to acquire when found. Therefore, outsourcing to a consulting company looks to be very appealing. But many of the real problems, the real failures in data warehousing, have occurred when the data warehouse has been outsourced.

The problem with using a large consulting firm to outsource the data warehouse development is that the consulting firm wants to treat data warehouse development like any other development project. The consulting firm wants a three-year contract where 10 to 20 consultants are brought in. The reason why the consulting firms want to treat data warehousing this way is because this is the way all their other contracts and projects are done.

The data warehouse should not be treated this way. Data warehousing should be built in small, fast iterations at a time. No single development effort in data warehousing should require more than three months. Therefore, from a cultural standpoint, data warehousing is a poor fit with the way that large consulting firms like to operate.

A related reason why large consulting firms do not do well with data warehousing is that they want to employ the develop-



ment methodology that they are familiar with. Unfortunately, in most cases, the large consulting firm is steeped in the tradition of the "waterfall" structured approach to development. In the waterfall approach, all requirements are gathered before the next step of development ensues. Each step of development is completed before the next step begins. This waterfall style of development is exactly the opposite of what the data warehouse requires.

Proper data warehouse development requires a spiral development approach to the building of the warehouse, not a classical "waterfall" approach. In a spiral approach a small subset of requirements are taken completely through the development process before the next set of requirements are tackled. Trying to get a large body of consultants to unlearn the knowledge and lore that they have carefully learned over the past 20 years is a delicate and difficult thing to do. For these reasons, when it comes to outsourcing, let the buyer beware.

HP Pro: How can an IT manager gain an accurate understanding of what data the company has, and what actually needs warehousing?

Inmon: The needs of what should be in the warehouse and how it should be

structured are determined only by the end user building multiple iterations of the warehouse. The end user cannot know what needs to be in the warehouse on the first and even the second iteration of development. The end user operates in what can be termed a "mode of discovery."

Only after the end user sees what the possibilities are, can they articulate what the real needs are. Therefore, the developer needs to quickly develop the warehouse in small, fast iterations so that the real end user requirements can be discovered. Trying to do a classical JAD session-based requirements gathering development, where all requirements are gathered before the first line of code is struck, is a prescription for disaster when it comes to data warehousing.


HP Pro: Do you have advice for an IT manager in a heterogeneous environment?

Inmon: It is normal for the legacy environment to exist in a heterogeneous state. The use of an ETL tool can greatly alleviate the work of integrating disparate data. As far as the different components of the warehouse existing in a heterogeneous state, that too is normal. The ODS is in one technology, the exploration warehouse is in another, and each data mart is in yet another technology. There is no reason why the different components of the corporate information factory should not exist in different technologies.

HP Pro: How can IT managers leverage their metadata? And how much metadata is enough metadata?

Inmon: Metadata is a tough topic. There exist few tools in the marketplace that are of any help. And what few tools there are center around the notion of a centralized repository, which is a mainframe idea that has little relevance to a distributed architecture.

One of the main reasons why there are no metadata distributed tools is that Sand Hill Road, where the venture capital community lives, refuses to fund any significant distributed metadata companies. Sand Hill Road has the attitude that they don't want to fund a company where money has never been made. Think about it. No company or metadata product has ever been financially successful. Therefore, the venture capital community doesn't want anything to do with metadata.



Bill Inman, "The Father of The Data Warehouse"

This leaves managers with a real dilemma. This is not a good answer for many reasons, but if you have to have metadata today, build your own solution, since the metadata products of the world that are suitable to data warehousing will not likely be forthcoming.

HP Pro: Is there a metadata standard?

Inmon: The only metadata standard I am aware of is that which is promoted by the Metadata Council.

HP Pro: Can you explain the concept of "cubes" or "cubing data?"

Inmon: Cubes refer to multidimensional data warehouses. Cubes are best placed in data marts, not large-scale, industrial-strength data warehouses. Cubes are the basis of OLAP technology.

HP Pro: What is today's IS manager missing out on (if anything) when it comes to data mining?


Inmon: The natural extension of data warehousing is data mining. Once the data has been collected, integrated and cleansed, it is natural to want to use the warehouse to start to find different patterns of transactions. Data mining placed on a data warehouse is a diamond placed on a golden ring.

HP Pro: What is the essential difference between a data mart and data warehouse?

The Role of SI in BI

Source: survey.com

In its BI/DW Research Group's Database Solutions III report, Survey.com recently asked companies to rate the importance of vendor-provided BI services (5 being the most important; 1 being the least important). For the complete results on BI/DW report, visit www.survey.com. The charts here and on page 25 indicate some of the more relevant findings.



Relative Importance of Vendor-Provided BI/DW Services (5 is most important)

Inmon: A data warehouse is fundamentally and architecturally different from a data mart. Comparing a data mart to a data warehouse is like comparing a tumbleweed to an elm tree. Both are plants, but both have an inherently different genetic structure.

A data warehouse contains a wealth of historical, granular, corporate and integrated data, and needs to have data structured in the most flexible manner possible.

A data mart contains a minimum of historical, aggregated and summarized data; is designed around a department's requirements; contains data optimized for departmental access; and is structured for optimal speed of access.

Because of these very stringent differences between data marts and data ware-

houses, the internal structure is quite different. Data warehouses are best normalized. Data marts are best denormalized into star joins and snowflake structures.

HP Pro: What's the most common DW mistake an IT manager faces?

Inmon: The single, largest mistake IT management makes is in trying to treat a data warehouse as they have treated online transaction processing systems. Development is radically different. Usage is radically different. Capacity planning is radically different. Users' attitudes are radically different. But IT managers think that the existing IT organization can build and operate a warehouse just as they have done for years with OLTP systems. In fact, the IT function must go back to square one. And people don't like to do that.

HP Pro: Is there a way to test the validity, accuracy and scope of a data warehouse?

Inmon: The validity, accuracy and scope of the warehouse is easily tested. Ask the end user to pay for the warehouse and you will find out more than you want to know.

HP Pro: What are some long-term maintenance issues facing a data warehouse?

Inmon: Far and away the largest long-term maintenance issue of the warehouse is that of managing the volume of data found in the warehouse. IT organizations are used to managing systems measured in megabytes and gigabytes. Data warehouses are measured in hundreds of gigabytes, terabytes and even petabytes. There simply is nothing to compare warehouse maintenance to, based on a background of OLTP.

The volumes of data present some novel challenges to the IT manager. The first of those issues is budget. Many a manager simply assumes that data in a warehouse should be stored on high performance storage. After all, that's the way we have always done it, and the hardware vendor seems to think this is a grand idea. But after about 500 giga-

bytes, see what happens. At some point in time there is the realization that "business as usual" will send your company to the poorhouse.

There comes a day that alternative forms of storage become the only rational way to succeed with a warehouse. Stated differently, those corporations building warehouses over a terabyte of data are throwing money away by the bushel by continuing to place their data on high-performance storage. And these are the very same people who complain about how much the data warehouse costs.

Simply stated, for the long-term vision of data warehousing, high-performance storage is absolutely not the appropriate technology.

HP Pro: What's the biggest challenge IT managers face when implementing a data warehouse?

Inmon: The IT manager faces both a managerial and developmental challenge in building the data warehouse. If I had to give advice to the IT manager, first of all, build the warehouse iteratively, one step at a time. It is poisonous to try to use waterfall development and the "big bang" approach to the building of the data warehouse.

Second, make sure you understand the framework known as the corporate information factory. There are good reasons why there are different components and dynamics.


Third, make sure the end user is involved from the beginning of the development of the warehouse. Without the end user's input, the warehouse stands the chance of becoming a technological masterpiece, but a business failure.

Be prepared for handling a volume of data that you have never seen before. Data warehouses surpass everyone's expectations when it comes to the volumes that accumulate in the warehouse.


Do not attempt to use the "waterfall" development approach to build the data warehouse. The warehouse requires a spiral development methodology.

Finally, check the references of your consultant. Do not accept your consulting firm's credentials as proof that your warehouse will be built properly. Instead, look at the individual resume of the people the consulting firm offers you, and examine - closely - their individual credentials.

For more information about data warehouse and the corporate information factory, visit Bill's Web site at www.billinmon.com. ♦



When a new DW is to be deployed in your organization, what type of vendor do you turn to first?




The Importance of Services Vendors to the Overall BI/DW Program (5 is most important)

ASAP Software Integrates E-Commerce Front End to Its HP e3000 Business Systems

ASAP Software, established in 1984, is one of the world's largest software resellers. The vendor sells and supports more than 40,000 PC software and add-on hardware products targeted primarily at large and mid-size organizations in business and government. In 1999, ASAP's Vice President of Systems, Bob Lewandowski, faced a real challenge.

Several years before, ASAP had developed a homegrown ordering package that allowed customers to place orders for software and services. In the interim, ASAP's business continued to grow – the firm regularly achieves a compound annual growth rate of 35 percent – and it needed a system with more scalability and user friendliness. It was time to replace the ordering package.

Lewandowski faced a number of issues. The homegrown package, written in C, ran on a front-end Windows NT system that interfaced with HP e3000s on the back end. The back-end applications – ERP and customer service modules – were written in COBOL. Lewandowski wanted, if possible, to select new software for the front end that would integrate with the COBOL applications on the HP e3000s. That meant avoiding software that had to be laboriously retrofitted to operate compatibly with the e3000s. Lewandowski also didn't want a "vanilla" solution that had to be reworked to synchronize with ASAP's primary business functions. As if all that weren't enough, the vice president had business relationships to consider. Microsoft is one of the company's largest vendors, and Lewandowski felt that he had to consider a Microsoft solution.



I spoke to Bob, I explained that OrderChannel is an ideal package to help companies like ASAP Web-enable their sales and supply chain and would be a good fit for them. Then I added that FRI also has the experience to do the integration with a COBOL application or the latest Microsoft suite."

"I knew Mark was familiar with the HP e3000, and that knowledge was something we needed," Lewandowski says. "I was well aware that there are a lot of e-commerce companies out there that can build

Web sites and storefronts. But we were not new to e-commerce and I did not want to duplicate business knowledge on a new platform."

Redwood used a product demonstration to convince Lewandowski and ASAP's senior management that OrderChannel would fit their short-term and strategic needs. FRI would provide a solution with OrderChannel running on Windows NT servers on the front end integrated with HP e3000s on the back end. Early on, the two companies agreed that all changes made in the COBOL on the HP e3000 side needed to be automatically reflected on the Internet, or NT, side.

The project took six months to complete and went into full production in May of 1999. The solution consists of a cluster of HP NetServers running the OrderChannel package. The databases are built using Microsoft's SQL Server. At the back end are HP

THE SOLUTION

Lewandowski had known Mark Redwood, President of Fioravanti-Redwood International (FRI), for several years. Both had served in administrative capacities in the Chicago Regional User Group (CRUG). Lewandowski was ready to listen when Redwood told him about an e-commerce solution that FRI had purchased in July of 1998 and tailored to integrate with the HP e3000. Called OrderChannel, the solution is designed to support online order and customer service functions. FRI had beefed up the software after the purchase, continuing to add business-to-business capabilities.

"Scalability and usability are the two biggest issues most companies face when it comes to e-commerce," Redwood says. "When

ASAP Technology Configuration

- HP e3000 987/150, MPE/iX 6.0
- HP e3000 969/400, MPE/ix 6.0
- HP NetServers (4 CPUs)
- Microsoft SQL Server 7
- OrderChannel 5.02

Feeling Lucky?



Taking a gamble with your IT operation? The odds of having an error-free data center are probably about the same as winning the lottery. So why play games of chance with important information.

We're willing to bet that JobRescue – our powerful error-detection and notification program – will put the odds in your favor. Whether you hit the lottery or use JobRescue, you'll come out a winner.

Just for contacting us, we'll send you a California Lottery Scratchers® ticket. No obligations.

nobix

Call 800.538.3818,
e-mail us at
sales@nobix.com or
check out our web site
www.nobix.com
Try your luck!

Nobix is a trademark of Nobix, Inc. Other products mentioned are trademarks of their respective owners. Limit one Scratchers® ticket per person. Any state laws concerning lottery winnings and any taxes associated with winnings are the sole responsibility of the individual. Offer expires 12/31/00.

**TOTAL SOFTWARE**

SOLUTIONS

**Specializing in Manufacturing &
Financial Software Since 1980.**

CFS Systems provide the tools you need to manage and control your business!

- ◆ Integrated Manufacturing
- ◆ Accounting
- ◆ Job Costing
- ◆ Distribution

Everything in CFS is stable, feature-rich, and easy to use and Y2K compliant.

**Contact CFS today
for a Free
CD-Rom Demo**

*Implementation, Installation, Training,
Modifications, Conversion Services & System
Consulting*

THE TOTAL SOLUTION

**11023 118th Place N.E.
Kirkland, WA 98034**

Phone (425) 803-3282
Fax (425) 803-6804

e3000 servers running distribution management systems, including both ERP and customer service modules. The data resides on multiple TurboIMAGE databases (see Figure 1).

An OrderChannel API links the NT servers to the HP e3000s through an online integration broker. Carry-forward shadow files speed processing on the NT cluster platforms while ensuring that all mission-critical information is centrally retained on the HP e3000s. These shadow files and database are automatically updated from the source databases in the event of a system failure. The same APIs allow programmatic changes to the COBOL programs in the back end to be carried forward to the NT boxes and reflected on the Web site.

“You would be amazed at how many e-commerce pages in use today have a total disconnect requiring the front-end



Bob Lewandowski,
Vice President
of Systems at
ASAP Software

orders to be collected and rekeyed into a separate back system system," Redwood says. "The ASAP system is integrated both in terms of data coming in from the Web agent and from the HP e3000 business backbone."

THE "E-WAY"

ASAP has branded OrderChannel to its customers as "E-Way." The package provides extranet connectivity so that business

and government users can order, manage and track technology, including software licenses, shrink-wrap, add-on hardware and technical services from the application.

OrderChannel allows ASAP's customers to access contract pricing, as well as track their orders. The newest features of the application enable it to automate numerous management and administration features, including setting user purchasing limits and order approval

routing. This allows companies to authorize spending limits throughout the organization, instead of requiring every single order to go through central purchasing. For maximum ease in order processing, customers have numerous searching and ordering options.

ASAP considers Order-Cash a resounding success. In the application's first year of operation, ASAP's online orders increased 340 percent. The size of orders also increased, as did the amount of online orders as a percentage of the total.

Visit ASAP Software's Web site at www.asapsoftware.com, or FRI's Web site at www.friinc.com. ♦

**Bill wants to hear
about your adventures,
victories, advice,
gripes and grievances
surrounding the HP e3000.
You can e-mail him at
pjkecom@earthlink.com.**

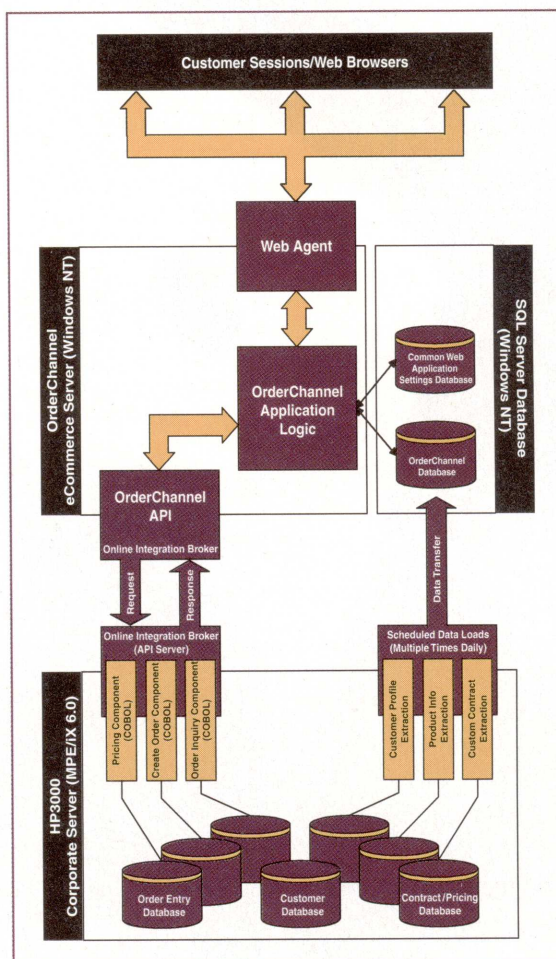
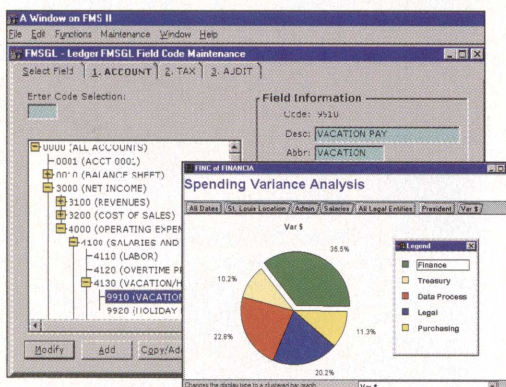
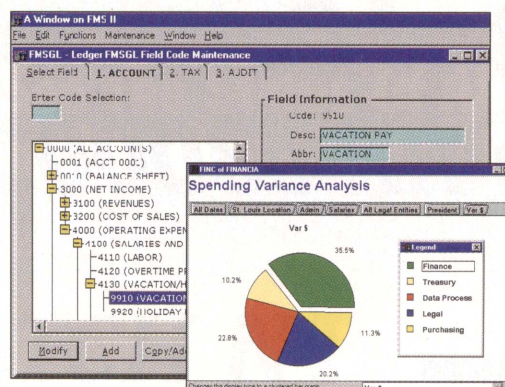


Figure 1: The ASAP Software Diagram

Client-Server Accounting Software for the HP e3000



FMS II on the HP e3000



FMS II on Windows NT

And Windows NT.
Same look. Same software.

Come see us at
HP World
Booth 319!

Fast, economical deployment

Superior functionality

Comprehensive training

Web enabled

Robust security

Our flagship product, FMS II, incorporates leading-edge technology and platform portability with the latest features in Windows design. So regardless of platform, the look and feel is exactly the same - Windows.

Plus, we've spent extra time so you won't have to spend extra money. If you change platforms, you don't pay for the software again. Nor do you pay for user retraining since what the user sees is exactly the same on both platforms.

So, whether you want your accounting system to run on an HP e3000 or if NT is in your future, FMS II will work for you.

**MITCHELL
HUMPHREY**
accounting software

800.237.0028

www.mitchellhumphrey.com

HEWLETT
PACKARD
Premier
Solution
Partner

Microsoft Certified
Solution Provider

Get the Most out of Data on the HP e3000

Jean Nattkemper

When HP's Commercial Systems Division (CSY) dropped an "e" in the HP 3000, rebranding it the HP e3000, CSY was telling the world that the platform is a player in e-commerce and e-services. That includes data warehousing and business intelligence.

"Many people don't know that," says Alvina Nishimoto, CSY's R&D Program Manager. "We hear people say, 'Oh, it's so hard to get the data I need out of my IMAGE database.' But it has nothing to do with the IMAGE database. It has to do with the fact that our applications have been very highly tuned for online transaction processing, and the data is organized as such. It's not well organized for reporting and the trending type of information people want from a business intelligence database."

In other words, production systems are organized to optimize online transaction processing, but business intelligence requires different types of organization (see Table 1). This is where tools come in. They're used to optimize data for reporting purposes.

"A lot of what customers want to do with data is perform queries that cannot be predicted ahead of time," Nishimoto says. "Take, for example, buying trends. Businesses want to be able to look at one thing, think that's interesting, and try a different query."

Well, the majority of databases have

to predefine these queries in order to get good performance out of their e-intelligence solution, so that kind of negates the benefit. That's why we emphasize tools – because customers are going to do ad hoc queries and do them on a large amount of data. And tools – not hardware – provide for this."

"This has nothing to do with IMAGE, specifically," Nishimoto explains. "Relational databases have similar issues with business intelligence and data warehouses, and so they require similar sorts of tools. Our tool vendor partners, like Taurus and DISC, sell their products for other kinds of databases as well."

CSY has had tool vendors as partners for some time because HP e3000 users have demanded business intelligence capa-

bilities. Take, for example, users of Smith-Gardner's mail order and cataloging system. "These users have been running WebOrder [Smith-Gardner's Web order fulfillment system] for quite a while," Nishimoto says. "They look for trends so they can get ideas, special promotions and so on that can increase sales."

DECISION VAULT: BI ON THE HP E3000

CSY offers a business intelligence solution for the HP e3000 that allows users to remain with the reliable and familiar IMAGE database and still obtain the



	Operational Database	Data Warehouse
Function	Data processing, support of business operations	Decision support
Data	Process-oriented, current values, highly detailed	Subject-oriented, current and historical values, summarized and sometimes detailed
Usage	Structured, repetitive	Ad hoc, some repetitive reports and structured applications
Processing	Data entry, batch, OLTP	End user initiated queries

Table 1: Operational Databases vs. Data Warehouses – Data organized for online transaction processing is not well organized for reporting and other business intelligence functions.

reporting and trending capabilities they need. The bundle, called Decision Vault, involves four HP partners: DISC, Taurus Software, Quest Software and Brio Technology (see Figure 1).

For the solution, Taurus and Quest jointly developed BridgeWare. BridgeWare performs an initial load of the data warehouse and carries out the essential data transformation, data scrubbing, data denormalization, and loading of data. BridgeWare updates through Quest's Netbase (SharePlex/iX), which detects changes automatically as they occur.

After the data has been loaded into the IMAGE warehouse, it is queried, analyzed, and retrieved by end users through DISC's OMNIDEX. OMNIDEX is a database search engine that supplements the existing database, supplanting the native database indexes with one specifically designed for high-performance, multidimensional access and dynamic data aggregations. It provides additional access not possible with the

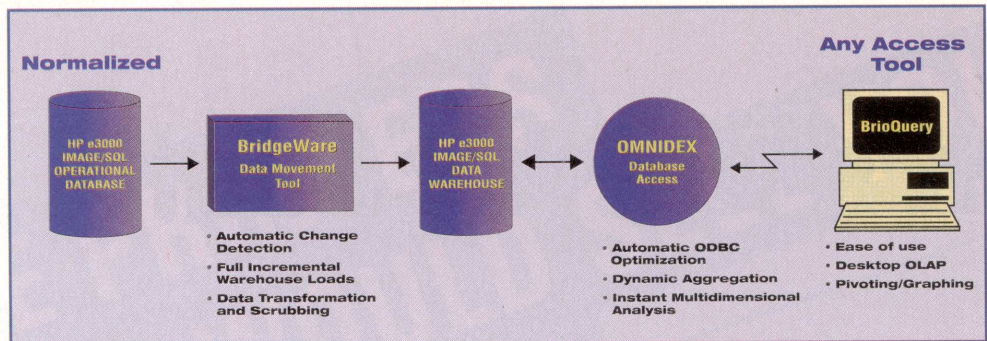


Figure 1: Decision Vault – BridgeWare provides data extraction services, OMNIDEX high-performance indexing, and BrioQuery, a query, analysis and reporting interface to the data warehouse.

native IMAGE database alone. DISC's indexing capability boasts impressive performance, as demonstrated by a test of an operational healthcare IMAGE database containing 9 million records and running on an HP 3000/967 in a batch-reporting scenario. In the test, complex queries were executed in 21 to 50 seconds.

OMNIDEX integrates transparently with the ODBC analysis/reporting tool BrioQuery. BrioQuery, designed for end users, provides an intuitive query,

analysis and reporting interface to the data warehouse.

Decision Vault provides a native business intelligence solution for the HP e3000. It is used especially by businesses in CSY's targeted vertical markets, including healthcare (Amisys), mail catalog (Smith-Gardner) and manufacturing (MANMAN). ♦

— Jean Nattkemper is the Editor at Large for HP Professional. She can be reached at jnattkemper@hpro.com.



Sales, Service, Support

- Hewlett Packard 3000, 9000, & NT Servers & Workstations, and all peripherals & upgrades.
- Our products are pre-tested and warrantied for up to 1 year.
- Consulting and troubleshooting via e-mail or telephone.
- Nationwide Onsite and Depot Service programs.
- 24-hour telephone support on hardware and operating systems.

Toll free 800-840-4277

Email: sales@blackriver.com

Fax: (440) 365-3916

www.blackriver.com

Performance tuning just got easier!



Understand exactly what's happening

SarCheck translates pages of sar and ps output into a plain English or HTML report, complete with system tuning and hardware upgrade recommendations.

Maintain full control

SarCheck fully explains each of its recommendations providing you with the information you need to make intelligent, informed decisions about your system.

Plan for future growth

SarCheck's Capacity Planning feature helps you to plan for growth, before slow downs or problems occur.

Available for most HP-UX, Solaris and SCO systems.

Visit us at <http://www.sarcheck.com/>

AURORA
SOFTWARE INC.

Tel 1-603-382-4200 Fax 1-603-382-4247



It's Not My Fault

The term “fault management” is an interesting one. In a book I read recently, it is defined as “detecting and reporting unusual or unacceptable behavior.” I suppose it is like the half-full/half-empty glass

syndrome. Would you rather manage system integrity or system faults? Does your job description say you are responsible for system faults? If so, what an easy job. You can spend the whole day going around unplugging network cables and peripherals, killing daemons, shaking disks and creating massive sparse files to overflow disks.

On the other hand, your job responsibility might include availability management. This would include monitoring the performance and behavior of the systems and networks under your control. You would be monitoring for abnormal conditions, so that you can fix them quickly if they arise, or before they affect availability.

Hmmm ... after thinking about it, maybe fault management is a pretty good term after all. This train of thought occurred recently when a book fell from heaven and landed in my snailmail box. The title, you could guess, is *UNIX Fault Management* (HP Professional Books Series, Prentice Hall) by Brad Stone and Julie Symons.

In the book, the authors make the assumption that the system operators are not the system administrators, which I found interesting. For many of us, they are the same person.

One of the first paragraphs in the book states that a study conducted by the GartnerGroup determined that operators are responsible for most unplanned downtime situations, which reinforces my thought that fault management really is a good term. It goes on to imply that automating system monitoring is a good idea, and one that I agree with. This leads

to what I think is one of the two best uses of this book.

If you are trying to make a decision as to which of the many products available to purchase for monitoring systems, read the book. It reviews usage, extensibility and capability, of several products, discussing both free and purchased methods. It also talks about using system commands for this purpose. I think this would be very helpful in making a decision, partly since the book specifically mentions HP-UX products.

The second great use of this book would be to make it required browsing for new system administrators. The second chapter defines major events in the life of a system, defining those pieces that should be recorded, tracked or detected. They are all good things for admins to be aware of, if they are not already.

Chapter three is a good overview of how monitoring tools might (and do) work, listing out the general behavior of several tools, such as IT/Operations, Unicenter ITG, SyMON, BMC Patrol, PLATINUM ProVision and Measureware. This is good reading, since it tries to inform you of the general approach each tool takes, and you can decide if that fits your perspective.

The remainder of the book goes into more details about each section of the system you might need to monitor (sys-

tem, disks, network, applications, databases and the enterprise as a whole). The book covers how the various tools perform, and, most important

to a cheap company like ours, using system tools to perform each area of monitoring.

One thing I found rather annoying about the book is that they did not use any visual divider when changing from discussions of one tool to another. I would strongly suggest something more visible than just a heading in a larger point size. How about a new page,


or a horizontal rule?

SOME FREEBIES

Here are some of the things I learned about using standard commands in just a brief browsing of chapter four.

The “file table full” error message often catches new system administrators. If you did not configure enough file table space, you can have a “full” disk, when most of it is empty. This can happen when not enough space is allowed for the “index” of files. I had never happened to read about the `-i` option to the `bdf` command before. When used, it adds two columns of output to the usual report. One is **iused**, the other is **ifree**. This lets you know quite easily if you are approaching a problem. In fact, a rather simple Perl program could be used to test for either type of disk full problem

See HP-UX on page 39



Fred Mallett
frederm@famece.com

Mining Continued from page 20

could be incorporated into the warehouse and the relative importance of incurring the cost of rebuilding the warehouse.

If the methodology and architecture used to implement the warehouse do not support a quick iteration cycle, then it might be easier to redesign the warehouse and move forward to capture the additional information required, to see if evidence can be gathered in the future to verify the hypothesis, rather than rebuilding the warehouse. The feasibility of either of these approaches is also a function of how quickly the warehouse and the programs that populate it can be modified, and how important it is to find a timely response to the business question raised by the anomaly uncovered through data mining.

An evolutionary architecture. As Agosta also points out, ideally one would like the results from data mining before designing what a data warehouse should contain, but this is not possible as such applications cannot be run meaningfully against raw operational data. For this reason, companies that want to avail themselves of data mining should ensure that the products and methodology they use in their warehouse implementation support a quick modification cycle. The key to this is keeping a metadata audit trail of everything discovered in building the warehouse – the source fields utilized and how they map to the fields in the warehouse; the look-up tables and business rules used to transform data values and how these are affected by schema changes encountered, etc.

“WIRED” FOR MINING

Data mining is appealing because it suggests how one can combine the computational muscle of the computer and techniques drawn from statistical analysis and optimization technologies with the creativity of the human mind to create a supersleuth. However, the drive toward e-commerce and what Giga calls the “zero latency enterprise” adds another wrinkle to the technical and conceptual complexities discussed above. With the goal of using the Web to support personalized marketing and just-in-time manufacturing, companies are using middleware products to implement near realtime warehouses. Depending on the business – for example, those involved with stock transactions – these near-realtime super-applications are mission-critical, but are they data warehouses?

Eleven Steps Continued from page 20

6. Leverage external resources. External information sources can increase the value of internal information. Rather than simply comparing yearly sales, use of external data might make it possible to compare sales growth against the increase in the overall market. The integration challenge is even greater when external data sources are involved. In some cases, external data will differ so drastically from existing schema that data transformation algorithms will be required to make use of the external resource.

7. Use information distribution methods. The biggest technological improvements in the data warehousing field have come in the area of information delivery. Today, users can subscribe to regular, personalized reports and have them delivered over e-mail, housing report data securely and economically on the server.

8. Focus on high payback marketing applications. Most of the hot applications in data warehousing involve marketing, because of the potential for an immediate payback in terms of increased revenues. For example, catalog manufacturers are using data warehousing to match specific customer characteristics and customer purchases to target their consumers' needs.

9. Emphasize early wins to build organization support. The availability of a wide range of off-the-shelf solutions has helped to reduce cost and lead-time requirements for data warehousing applications. Off-the-shelf solutions won't complete project objectives, but they can be used to provide short-term point solutions that serve as a training and demonstration platform.

10. Don't underestimate hardware requirements. The hardware requirements for data warehousing database servers are high, due to the large number of CPU cycles required to slice and dice data. Database size also plays a part in server performance requirements. Be sure to select a scalable platform; the typical data warehouse implementation starts out at the departmental level and grows to an enterprisewide solution. Purchasing servers that can be expanded with additional processors is one possible approach. A more ambitious idea is to combine loosely coupled systems that enable the database to be spread out over multiple servers, appearing as a single entity.

11. Consider outsourcing development and maintenance. Most data warehousing applications fit the criteria for a good outsourcing project – a large project that has been defined to the point that it does not require daily interaction between business and development teams. There have been many cases where an outsourcing team is able to make dramatic improvements in a new or existing data warehouse. Typically, these improvements do not stem from an increased level of skill on the part of the outsourcing team, but rather flow from the nature of outsourcing. The outsourcing team brings fresh ideas and perspective to their assignment and can use methods and solutions developed on previous projects.

– Phillip Blackwood is the Vice President of Data Warehouse and Business Intelligence at Syntel Inc. (Troy, Mich.; www.syntelinc.com).

Historically, data warehouses were intended to assist management in introspection. The concept of introspection suggests some element of elapsed time; the notion of time outside the turmoil of day-to-day business that allows us to devise a better plan. Even without this time conflict, the question remains whether even the best combination of data warehousing and data mining can anticipate a paradigm shift where many of our essential assumptions are no longer

relevant. The need to balance our ever increasing need for speed, efficiency and innovation with wisdom and an understanding of our limitations is surely our greatest challenge. ♦

– Katherine Hammer
is President and
CEO of Evolutionary
Technologies
International
(Austin, Texas).



Updating Your Microsoft Certification

With the arrival of Windows 2000, it's time to update your professional certifications from Microsoft. The release of Windows 2000 has offered Microsoft an opportunity to update the very

popular Microsoft Certified Systems Engineer (MCSE) program and they are taking it. They appear to be responding to various anecdotal stories about so-called "paper MCSEs," certified because they passed the tests, not because they are capable. Now, Microsoft is making it harder to become an MCSE.

While, in the short term, this is annoying and expensive for those of us who must update their certifications, in the long term it's good for everyone. For Microsoft, it helps the MCSE title retain value and have actual meaning in the real world. For employers and potential employers, it provides added assurance that candidates who are MCSE will have the technical skills necessary to work effectively. Microsoft seems to have aimed the new requirements squarely at weeding out the paper folks and providing some assurance to potential employers that the MCSE will be a real certification with meaning. For those folks seeking certification, it provides more credibility when seeking employment.

So what has Microsoft done? Aside from revamping the test requirements, they are aggressively retiring exams. In the past, certification exams were usually good for at least two versions. In other words, your NT 4.0 exam and certification would be valid until the Windows 2000 (or whatever) certification was released. Now, Microsoft is retiring the NT 4.0 exams, effective December 31, 2000. People who became certified with 4.0 exams will retain their certification until December 31, 2001. Microsoft states that this is to help "organizations ensure that they have professionals on board who are fully skilled and ready to use the latest technology." This accelerated retirement and totally redesigned certification track will probably force existing MCSEs to start taking tests as soon as possible.

NEW MCSE REQUIREMENTS

Previously, there were two MCSE tracks: NT 3.5x and NT 4.0. These consisted of six tests, including four core tests related to the operating systems, clients and networking basics and two elective tests, allowing the MCSE to pick an area of concentration.


The Windows 2000 certification requires seven tests, including five core exams and two electives. Microsoft chooses four of the core exams. The MCSE candidate can choose the fifth from three options. The core requirement exams are:

- Installing, Configuring and Administering Windows 2000 Professional (Exam 70-210)
- Installing, Configuring and Administering Windows 2000 Server (Exam 70-215)
- Implementing and Administering Microsoft Windows 2000 Network Infrastructure (Exam 70-216)
- Implementing and Administering Microsoft Windows 2000 Directory (Exam 70-217)

The fifth core exam is chosen from:

- Designing a Microsoft Windows 2000 Directory Services Infrastructure (Exam 70-219)
- Designing Security for a Microsoft Windows 2000 Network (Exam 70-220)
- Designing a Microsoft Windows 2000 Network Infrastructure (Exam 70-221)

This optional core seems to be designed to let the MCSE specialize in a particular area. The basic Win 2000 Server exam will certainly discuss security; but, by taking the Designing Security exam, an MCSE could say they specialize in security. This may have value for particular employers.



Ryan Maley
ryan@maley.org

The elective exams are chosen from the same pool of exams as previous tracks. A wide variety of topics are available including Proxy Server, Internet Information Server, SQL Server administration, SQL database design, Systems Management Server and Site Server. One new exam, Upgrading from NT 4.0 to Windows

2000, has been added. Additionally, the optional core "Designing" exams can also be used towards the elective requirement.

ANY QUESTIONS?

Microsoft has also announced some changes to the exam questions. In the past, questions were primarily multiple choice. This, theoretically, allowed the candidate to guess the correct answers and pass the exam. Occasionally,

a question would require multiple answers or a graphical point to a particular area on a graphic. With the new exams come two new question types: case study and select-and-place. Additionally, Microsoft will probably make more use of simulation questions that have started to appear in some exams.

The case study questions present a scenario and offer various solutions. They are designed to force the candidate to actually solve a problem.

Contrast this with simply knowing a fact. For instance, you may know that the length limit of various cable types. Does this mean you can troubleshoot a network? Cable lengths are simply one aspect of network design and troubleshooting and probably not the most important one.

Select-and-place questions are just that. They will require the candidate to


manipulate graphics or text on the screen. For instance, drag a Web server to the correct side of a firewall for full access. The simulation questions, cursed by many exam takers, will force you to perform actual tasks during the exam through a simulated interface. Instead of asking, "How do you change the bindings on a network card?" and presenting the multiple answers, a candidate would have to go to the Control Panel, select Network, etc.

By moving away from memorized facts to more "real-world" problems, the candidate is required, by default, to have some experience. Indeed, one of Microsoft's stated goals for the new exams is to demand "the candidate have troubleshooting skills acquired through hands-on experience and working knowledge." Microsoft expects that candidates will need a least one year of hands-on experience to pass the exams.

Speaking as an MCSE, the whole thing is annoying. I'll have to spend a lot of time getting up to speed – not to mention spending \$100 to take each test. However, in the long term, the changes are good. Win 2000 MCSE will add more

credibility to their certification and, hopefully, we'll get rid of some paper MCSEs who don't know seem to know coax from twisted pair. ♦

Resources



Microsoft Certification:
www.microsoft.com/mcp

Microsoft Training:
www.microsoft.com/train_cert


Saluki MCSE Mailing List:
www.saluki.com/maillist.htm

An excellent resource for people who self-study and want to keep up on certification.

Prometric: www.prometric.com;
and Vue: www.vue.com

**Want to take an exam or find out what's involved?
These are the two official test providers.**

— Ryan Maley is a Microsoft Certified Systems Engineer who needs to work on the 2000 track. He can be reached at ryan@maley.org.



**BUY • SELL
LEASE**

**HIGHEST QUALITY
REFURBISHED
HP EQUIPMENT
WITH PROFESSIONAL SUPPORT
AT A
FRACTION
OF THE HP PRICE**

- 9000/3000 Equipment
- Immediate Availability
- Up to 80% Off HP List Price

NHCSI
New Hampshire
Computer Systems
International LLC

877-HPNHCSI toll free
www.nhcsi.com

Your Competitive Source for HP Systems & Peripherals

Twin Cities Digital

HP 9000/SUN Remarketing Specialist

Top Quality Equipment

Excellent Pricing

888-880-0522

www.tcdigital.com

Fax 612-895-0521

sales@tcdigital.com

BUY

SELL


LEASE

RENT

E-mail: iodata@mindspring.com

HP 3000
HARDWARE

HP 9000
HARDWARE



also featuring

Wearable Computers

Sun, Digital, Cisco, Ascend

I/O DATA SYSTEMS, INC.

27378 W. OVIATT BAY VILLAGE, OH 44140

(440) 835-2211 FAX (440) 835-0220

<http://www.iodatasys.com>

GENISYS

CORPORATION

HP Hardware Solutions / Compatibles
Buy / Sell / Lease / Rent


GENISYS CORPORATION

14950 NE 95th ST. SUITE D, REDMOND, WA. 98052

phone 425 869 6055 • fax 425 881 1939

www.genisyscorp.com

DATA NEVER SLEEPS...



Be Seen Here

For

Reseller, Rental, Recruitment
and Training Opportunities

Call or email

Mary Thomas

972-664-6564

mthomas@101com.com

HP NETSERVERS – 9000

IBM NETFINITY RS/6000 • COMPAQ
PROLIANT • SUN MICRO

Systems, Parts,
Peripherals, Service,
E-Commerce Development
New & Refurbished

1-888-696-3344

ideal
SYSTEM SOLUTIONS, INC.

10200 - 73rd Avenue N, Suite #122
Maple Grove, MN 55369

Ph: (612) 425-8163

Fx: (612) 425-4931

Check our website:

www.idealssi.com

for more product and
service information

Ideal is an (8A) Certified Woman-Minority Owned
and Small Disadvantaged Business

Hard Disk & Floppy Disk Drives

*ISA Model 2085
(HP9122C Replacement)

*ISA Model 1015
(HP9153C Replacement)



Saaya Inc.

Tel: 281-493-9925

Fax: 281-493-9926

e-mail: saaya@ev1.net

14760 Memorial Drive
Suite 105
Houston, TX 77079

HP CAREER SPECIALISTS

INFOTECH SEARCH GROUP, LTD.

"Professionals serving Professionals"

- We specialize in the placement of **PERMANENT** and **CONTRACTING** Hewlett-Packard talent.
- Let experienced professionals with a proven track record of success **work for you.**

Call Brian M. Donohue

Phone: 910-454-0132


Fax: 910-454-9547

P.O. Box 10010

Southport, NC 28461

E-mail: infotech@bcinet.net

SPATCH



TIRED OF BABYSITTING YOUR SYSTEMS?

Use Spatch instead.

Spatch can be configured to alert your pager or phone with full text every time your system hiccups.

WITH SPATCH YOU CAN RELAX!

Starting at only \$199

770-840-7377

DOWNLOAD A FREE DEMO AT WWW.SPATCH.COM

Hewlett Packard Systems & Peripherals 9000/3000 & Netserver

Unix Workstations • Unix Servers

- *Disk Drives
- *Monitors
- *Upgrades
- *Memory

**Full Warranty & Guaranteed
Eligible for HP Support**

877-667-8383

Buy/Sell/Rent

On Queue Computer



"Professional Service & Integrity"

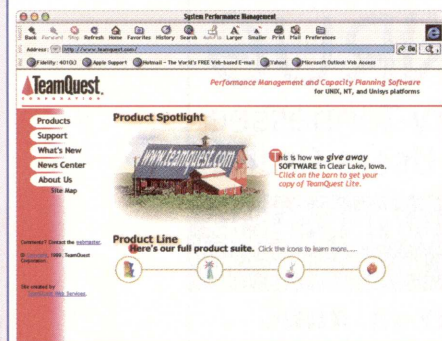
2 Martel Way, Georgetown MA, 01833

FAX (978) 352-2828

E-Mail Address Steveq@On-Queue.com

web.browser

www.teamquest.com



Improve performance and extend the life of your HP Systems. Come direct to the source.

**Contact your sales representative for Web
Browser marketing opportunities.**

WWW.HPPRO.COM/SUBSCRIBE

TYPE

YOUR INFORMATION.

CLICK

THE APPROPRIATE BOXES.

SUBMIT

YOUR ORDER TODAY!

**SUBSCRIBING TO THE INDEPENDENT GUIDE
FOR INTEGRATING HEWLETT-PACKARD
SYSTEMS IN YOUR ENTERPRISE COULDN'T BE
EASIER. AND IT JUST TAKES A MINUTE. TYPE,
CLICK AND SUBMIT. START YOUR FREE
SUBSCRIPTION TO HP PROFESSIONAL TODAY!**

(not enough space, or not enough index area). The program lines:

```
open(BDF,"bdf -i |");
junk=<BDF>;
while(<BDF>){
  @data=split;
  $cap=$data[5] / $data[6] * 100;
  print "$data[0] $cap\n";
}
```

Might result in something like:

```
/dev/dsk/c0t5d0 89.80319
/dev/dsk/c0t6d0 23.304
```

This is an indicator that the first disk is getting into trouble (89 percent of the allocated inode space is used up). If you then compared disk free space with this, we might find that this is normal (full disk and full inode space is okay). If the disk free space is high, but we are running out of inode space, this would be very bad.

Another surprise to me was the **ipcs** command. I had been asked in the past about a way to list shared memory segments. The book mentioned that there was such a command. It also lists active message queues, and semaphores. Another

memory jogger for me was the **sysdef** command which lists kernel configurations.

Most of the usual system monitoring commands are mentioned, though in no real detail, or usage. That is why I recom-

found no surprises there, but the section on available monitoring programs was quite interesting, since I have little experience with them.

I learned a lot by reading through

descriptions of the different tools, and would want to read this if I was responsible for making a decision about an impending purchase.

The book also has a couple of case studies that give examples of detecting and correcting problems with memory and disk mirroring which would be

The "file table full" error message often catches new system administrators. If you did not configure enough file table space, you can have a "full" disk, when most of it is empty.

mend that new admins browse the book to become familiar with available methods. Some of the commands mentioned for system monitoring are mailstat, ps, sar, iostat, ioscan, swapinfo, top, etc.

The disk and network chapters do a good job of listing out available standard commands and what they are used for. I

good for newer administrators to get a taste of real life.

Chapter eight of the book is about databases, which is an area I have little experience with. I think I'll go read it again. ♦

And if you agree that it's not Fred's fault, he can be reached at frederm@famece.com.

Advertising Sales Offices

(215) 643-8000
Fax (215) 643-3901

Midwest/East (215) 643-8063

Daniel J. Ferro, National Sales Manager
1300 Virginia Drive, Suite 400
Fort Washington, PA 19034
FAX (215) 643-3901
dferro@101com.com

West Coast (714) 628-0757

Carolyn Aliotta, Regional Sales Manager
7229 E. Clydesdale Avenue
Orange, CA 92869
FAX (714) 628-0758
caliotta@101com.com

Europe (410) 897-0297

Fran Grega, European Sales Manager
523 Samuel Chase Way
Annapolis, MD 21401
FAX (410) 897-0298
frangrega@compuserve.com

List Rental Manager (215) 643-8047

Cathy Dodies
dodiescl@101com.com

Marketplace Sales Manager

(972) 664-6564
Mary Thomas
FAX (972) 669-9909
mthomas@101com.com

Interactive Advertiser's Index

Marketplace

Aurora Software	31	Black River Computers	37
Black River Computers	31	Genisys Corporation	36
CFS	28	Hyde Company	38
Enterprise Group Ltd.	19	IDUG	37
Forsythe Technology	3	I/O Data Systems	36
Hewlett Packard	C2	Ideal Systems Solutions	36
Hillary Software	11, 13	Info Tech Search Group, Ltd.	36
IBM	15	Lynne Company	37
MB Foster	C4	Monterey Bay Communications	37
Mitchell Humphrey	29	On Queue Computers	38
Nobix, Inc.	27	Saaya Inc.	36
NHCSI	35	Surety Systems	37
Stratus	5	Ted Dasher & Associates	37
Sungard	C3	Twin Cities Digital	36
Technical Scientific App	2		
Xerox	9		

INFRATOOLS NETWORK DISCOVERY

Peregrine Systems' InfraTools Network Discovery is an automated network management system. Information collected by InfraTools Network Discovery is presented in a port-to-port, connection-level, graphical map of the network, which allows users to see where faults may occur, the potential impact, and the best way to resolve problems and deal with change. DNS name, IP address, MAC or SNMP field automatically identify every networked device over all media types, including virtual circuits. The system also monitors usage and traffic on the network and provides detailed statistical data on all devices, ports and lines. All information provided is hyperlinked to allow users to navigate between reports. *For more information, visit www.peregrine.com.*

STAT 3.0 SECURITY SYSTEM

Sunbelt Software introduces the STAT 3.0 (Security Test and Analysis Tool), a security system that is manned by a team of anti-hacking specialists. STAT 3.0 is designed to perform a complete security vulnerability analysis of Windows NT network resources and to provide updates through Web downloads as security breaches are identified and remedies are determined. STAT 3.0 features greater scalability and flexibility of network configuration enhanced screen navigation, detection of hidden machines by typing the IP address and increased testing of fixes. *For more information, call (800) 336-3166, or visit www.sunbelt-software.com.*

LEVEL 8, HP E-SERVICES

Level 8 Systems and HP have formed a marketing alliance to deliver e-services applications for the new Internet-capable HP e3000 platform. HP's newly renamed HP e3000 (formerly the HP 3000) provides the Web-enabling technologies required to develop back-office applications for the Internet. HP and Level 8 Systems plan to co-market Geneva Message Queuing to HP 3000 customers. Level 8 Systems' Geneva Message Queuing middleware software is an implementation of Microsoft Message Queue Services (MSMQ) for a full range of non-Windows operating systems. Geneva Message Queuing allows application developers to convert HP e3000 business applications to Web-based applications to support an e-services model. *For more information, visit www.level8.com, or www.hp.com.*

SANDISK COMPACTFLASH CARDS

SanDisk Corporation will supply HP with 16MB CompactFlash (CF) memory cards to be included with each new HP PhotoSmart C912 and HP PhotoSmart C618 digital camera sold. The CF cards will serve as the digital film for HP's new cameras, store high-resolution images and support state capabilities. SanDisk CompactFlash memory cards are currently available in capacities ranging from 4MB to 128MB. *For more information, call (408) 542-0500, or visit www.sandisk.com.*

FABRICREADY PARTNER PROGRAM

McDATA Corp. introduced a new level of support for implementing storage area networks (SANs) in an open system environment with its FabricReady Partner program. Strategic partners who participate in the FabricReady Partner program will be able to leverage the hardware and software applications testing done to ensure the completeness of their own evaluations. McDATA includes the major components of a SAN, from host bus adapters to storage systems and key software applications. *For more information, visit www.mcddata.com.*

HP VISUALIZE WORKSTATIONS

HP introduces the HP VISUALIZE C3600 and J5600 workstations. The advanced technical-computing workstations are based on the PA-8600 RISC processor and offer a 38 percent performance increase over HP's PA-8500 RISC processor. The HP VISUALIZE C3600 and J5600 workstations feature up to 8GB of SDRAM and a total on-chip cache of 1.5MB for reduced latency and increased performance. Both workstations support the HP VISUALIZE fxe, fx4 Pro and fx6 Pro graphics subsystems.

HP has also introduced the HP VISUALIZE B2000 workstation for UNIX, which features the HP VISUALIZE fxe full-featured graphics subsystem. The HP VISUALIZE B2000 system provides engineers and designers in Electronic Design Automation (EDA) and Mechanical Design Automation (MDA) with a balance of graphics, power, stability, price and computing performance in a UNIX system. *For more information, visit www.hp.com.*

HP, TROY EXTEND JETSEND PROTOCOL

HP and TROY Group announced TROY Group's NetSend line of Internet print

servers, which allow a wide range of non-HP and HP printers to communicate via HP JetSend. JetSend is an open communications protocol developed by HP that allows two devices to communicate directly with each other, without a PC or software drivers. *For more information, visit www.hp.com, or www.troygroup.com.*

HAMMER SLPRO12-FC/FC-FT


Hammer Storage offers the SLPRO12-FC/FC-FT Fibre Channel RAID storage subsystem for digital media applications. The SLPRO12-FC/FC-FT is a fault-tolerant, 19-inch, rack mountable, desktop or tower Fibre Channel storage subsystem. It supports up to 12 FCAL disk drives per enclosure and has a storage capacity of 600GB per enclosure, expandable up to 25 TBs.

The SLPRO12-FC/FC-FT also includes the JackHammer host bus adapter and a Fibre Channel to Fibre Channel RAID controller. The SLPRO12-FC/FC-FT supports RAID levels 0,0+1, 1, 3, 5, 10, 30 and 50 with online capacity expansion and an easy-to-use configuration utility. *For more information, call (510) 608-4000, or visit www.hammerstorage.com.*

POWERTERM PRO VERSION 8.0

Ericom Software's PowerTerm Pro version 8.0 is designed to incorporate Microsoft's Visual Basic for Applications (VBA). PowerTerm Pro 8.0 maximizes enterprisewide productivity by enabling access to accounting, inventory, management, transaction processing and other mission-critical legacy systems. It allows users to move data from a terminal-legacy document into a PC-based program using the VBA technology.

PowerTerm Pro 8.0 runs on Windows 95/98/NT and Windows 2000 and supports IBM Mainframe, IBM AS/400, UNIX, DEC VAX/VMS, SCO and Data General. *For more information, call (201) 525-5511, or visit www.ericom.com.* ♦



WE'VE GOT JUST ONE SMALL DETAIL FOR YOU TO CONSIDER WHEN YOU PLAN FOR BUSINESS CONTINUITY.

WHAT HAPPENS WHEN PLAN B FAILS?



It could be a fire. The worst storm in decades. Or a deadly virus. But when it hits, it could knock out your system for hours, even days. And you know how much that could cost. Unless you have the right business continuity plan in place. SunGard can help, with a range of highly responsive, surprisingly affordable options. (You will even spend less on a pre-configured SunGard recovery package than you would to buy a backup server.) And now we apply our disaster recovery experience and aggressive pricing to your other critical business needs. With SunGard's Internet and high-availability services, the people who need it always have access to your company's information. Call us or visit us online now. **YOU'LL GET OUR POPULAR WEB SITE GUIDE, e-PERFORMANCE, FREE.** Then relax. When plan B fails, SunGard won't.

SUNGARD® SECURE
THE NET BENEATH YOU.

1-800-813-7233 EXT. 130 OR
WWW.SUNGARDPLANB.COM/130

SUNGARD®

For over 20 years, M.B. Foster has been forging through the HP3000 marketplace to bring you complete data access and delivery solutions. Whether your need is cross-platform accessibility, client server technology beyond ODBC, multi-thread capabilities or business-to-business electronic commerce, M.B. Foster can handle all your data access requirements.



Win \$20,000 US
"Webify Your HP 3000" Contest!
Visit www.mbfoster.com for details

FORGING THE FUTURE OF SOFTWARE INTEGRATION

DataExpress is the definitive tool for accessing and reporting on HP3000 systems. But, DataExpress goes beyond reporting. It provides fast, reliable and direct access to your data whether it sits on a desktop, web browser, HP3000, HP9000 or Windows NT. DataExpress will integrate, extract and present your data in the format you want.

For EDI/EC, M.B. Foster offers two important families of integrated solutions - EDIWindows™ and AMTriX™ Messaging System. EDIWindows providing Job Control, Translation and Mapping is available for HP3000, HP9000 and Windows NT platforms. AMTriX enables Intelligent Messaging, application messaging, EDI and enterprise-wide electronic commerce for HP/UX and Windows NT platforms among others.

Do you have challenges in accessing and integrating your data? We can help. For more information about the latest technologies that will enhance your existing infrastructure, call M.B. Foster at 1-800-ANSWERS (800-267-9377) or visit our website at www.mbfoster.com.

M.B. Foster Associates

